

The Shadow of Uncertainty: Climate Policy and the Value of Petroleum Resources

BY Ingrid Emilie Flessum Ringstad, Giacomo Benini,
Valerio Dotti and Kyriaki Tselika

DISCUSSION PAPER

NHH



Institutt for foretaksøkonomi
Department of Business and Management Science

FOR 22/2025

ISSN: 2387-3000

July 2025

The Shadow of Uncertainty: Climate Policy and the Value of Petroleum Resources*

Ingrid Emilie Flessum Ringstad¹ Giacomo Benini¹ Valerio Dotti²
Kyriaki Tselika¹

July 7, 2025

¹Department of Business & Management Science, Norwegian School of Economics

²Department of Economics, Ca' Foscari University of Venice

*Preliminary draft. Please do not cite or circulate.

Abstract

We study how uncertainty about future climate policy affects the valuation of oil resources. Using a structural model of extraction and exploration applied to field-level data from the Norwegian Continental Shelf, we estimate the impact of climate policy uncertainty (CPU) on the shadow prices of both discovered and undiscovered oil. We find that higher CPU lowers these marginal values, especially after the 2015 Paris Agreement, reducing incentives to extract and explore. This decline translates into an implicit carbon cost of \$15–\$38 per tonne of carbon dioxide emitted. Unlike a Pigouvian tax, this shadow cost does not scale with emissions intensity or generate fiscal revenue. As a result, it reduces production and emissions in a diffuse and economically inefficient manner, without rewarding low emitters or financing green transition policies.

JEL Classification: Q35, Q58, C51, D81

Keywords: Oil Industry, Climate Policy, Decision-Making under Risk and Uncertainty

1 Introduction

In recent years, ambiguity has become a defining feature of Norwegian climate policy. At COP28 in December 2023, Prime Minister Jonas Gahr Støre joined international leaders in endorsing “a collective commitment to phase out the use of unabated fossil fuels” and emphasized that “phasing out means also rapidly phasing in renewables.”¹ Just a few months later, in the summer of 2024, Minister of Climate and Environment Andreas Bjelland Eriksen publicly stated that he did not foresee a specific date for phasing out oil production on the Norwegian Continental Shelf.² The contrast between these two statements reflects a broader tension, which characterizes many hydrocarbon-rich economies. While there is widespread agreement on the need to address climate change, there remains considerable uncertainty about how and when major fossil fuel producers will implement concrete measures to transition away from oil and gas (Stern, 2006; Weitzman, 2007; Nordhaus, 2007).

This ambiguity is not merely rhetorical; it translates into real uncertainty — particularly regarding the timing, scope, and credibility of climate policies — that plays a central role in shaping production and investment decisions in the oil and gas sector. A key source of this uncertainty is the erosion of policy narratives — shared storylines that anchor long-term expectations (Shiller, 2020), such as “we are phasing out oil and gas.” When such narratives are contradicted or diluted by short-term statements or mixed signals from policymakers, credibility suffers, and firms face growing difficulty in forming stable expectations about the regulatory environment.

Climate-policy uncertainty is a qualitatively different challenge from the well-studied problem of fiscal-regime instability in natural-resource markets. Volatile fiscal instruments—royalty rates, profit taxes, cost-recovery rules—are known to distort extraction incentives and discourage long-term investment, which is why countries such as Norway have cultivated a reputation for predictable petroleum taxation that balance investor confi-

¹See transcript of Prime Minister Støre’s address at COP28.

²See [Dagbladet \(2024\)](#) (in Norwegian).

dence with public-revenue goals. Climate targets and transition pathways, by contrast, are forged through political negotiation, framed in loosely defined language, and embedded in dynamic public discourse; they are not always codified in law and can shift with election cycles or external pressure. Consequently, even firms operating under stable fiscal frameworks now face growing uncertainty over whether, when, and how governments will impose binding constraints on fossil-fuel production and exploration to meet decarbonization objectives (Basaglia et al., 2021; Berestycki et al., 2022).

To investigate how this emerging form of uncertainty affects the petroleum sector, we build on the extraction–exploration framework originally proposed by Pesaran (1990), extending it to incorporate environmental regulatory uncertainty as a forward-looking cost shifter. In our model, each oilfield faces heterogeneous revenues and costs. Marginal revenues are estimated using a field-level pricing equation that isolates global demand-driven fluctuations from crude-specific price differentials. Specifically, we regress the spread between twenty-two traded oil classes and the Brent benchmark — the reference price for Norwegian oil producers — on a quality premium (or discount) that captures how much lighter and sweeter each crude is relative to the global average, over the period 1995 to 2023. This procedure yields estimates of marginal revenue for more than 120 oilfields³ located on the Norwegian Continental Shelf. Marginal costs are modeled as the sum of production costs and regulatory compliance costs, with the latter depending on both the *level* and the *variance* of national climate policy. To estimate the cost function, we merge field-level data on production and reserves from the Rystad Ucube database with a newly constructed Climate Policy Uncertainty (CPU) index that captures fluctuations in the regulatory environment over time. Specifically, we use this index as a proxy for the degree of uncertainty firms face about future climate policy and its potential impact on operational costs and investment decisions.

The CPU index is constructed using a media-based methodology originally developed

³This is based on the available data on oilfields found in the Rystad Energy Upstream Solution Ucube database

by Baker et al. (2016) and later extended by the OECD (Berestycki et al., 2022), who pioneered the use of newspaper text analysis to quantify policy-related uncertainty. We adapt this approach to the Norwegian context by developing a tailored lexicon of more than forty Norwegian search terms, grouped into three semantic categories: *climate*, *policy*, and *uncertainty*. These terms were translated from the OECD’s English-language vocabulary and refined through manual review to reflect Norwegian linguistic nuances and the structure of national policy discourse. The index is constructed from article counts retrieved from three widely circulated Norwegian newspapers — *VG*, *Aftenposten*, and *Dagens Næringsliv* — accessed via the Retriever media archive. For each newspaper, we compute the ratio of articles containing at least one term from each of the three semantic categories to the total number of articles published in a given period. These ratios are then standardized to account for differences in publication volume and editorial practices across sources. Following the OECD’s methodology, we average the standardized series across newspapers and normalize the result to generate a consistent time series that captures the evolution of climate policy uncertainty in Norway.

To assess the welfare implications of the CPU, we combine field-level estimates of marginal revenues, extraction costs, and discovery rates to recover the shadow prices of both discovered and undiscovered oil. These shadow prices reflect the opportunity cost of extracting an additional barrel — whether from existing reserves or through future exploration — and both are shaped by the level of CPU. The shadow price of discovered oil is directly influenced by CPU, as regulatory uncertainty enters the marginal cost function via compliance-related expenditures. In turn, the shadow price of undiscovered oil is derived from the value of discovered oil, scaled by the marginal return on exploration investment and adjusted for intertemporal discounting.

Having estimated the shadow prices of both discovered and undiscovered oil, we construct a counterfactual scenario in which the CPU is fixed at its minimum observed level. Comparing actual and counterfactual valuations allows us to quantify the deadweight loss

attributable to regulatory uncertainty. On average, the shadow price of a discovered barrel declines from 69.09 USD (median: 63.34 USD) to 68.39 USD (median: 61.39 USD), corresponding to a mean reduction of approximately 0.70 USD per barrel. In the post-Paris period (2015–2022), this wedge widens to more than 1.3 USD per barrel, reflecting the sharp increase in regulatory ambiguity following the international climate commitments of 2015. Since the shadow price captures the marginal value of extraction, these reductions translate into concrete fiscal costs for the Norwegian state. By multiplying the per-barrel shadow price reduction by each year’s aggregate oil production, we compute the annual revenue loss. Over the full sample, the average oilfield has lost 6.53 million USD per year due to increased uncertainty. This figure rises to 11.01 million USD annually in the post-2015 period — an increase of 68.6%. To contextualize these fiscal losses in environmental terms, we divide them by the aggregate upstream emissions obtained using the Oil Production Greenhouse Gas Emissions Estimator (OPGEE) obtaining an implicit “shadow carbon tax” — the marginal cost of regulatory uncertainty per tonne of carbon dioxide equivalent emitted. Across the full sample, the average implicit tax is 15.4 USD per tonne of carbon dioxide equivalent emitted. In the post-2015 period, this rises sharply to 38.4 USD per tonne, with the lower bound increasing from zero to nearly 20 USD and the upper bound holding steady at around 62.9 USD. These magnitudes rival those of formal emissions trading schemes, suggesting that climate policy uncertainty has introduced *a de facto carbon price* — one that is uncoordinated, unintended, and highly distortionary.

This reveals the central insight of our analysis: uncertainty about future climate policy functions as a shadow tax on fossil fuel extraction — one that is uncoordinated, non-transparent, and economically inefficient. Unlike a Pigouvian tax, which scales with emissions intensity and generates fiscal revenue, this implicit cost affects high- and low-emitting producers alike, without rewarding cleaner operations or providing funding for green transition efforts. The result is a diffuse reduction in production and exploration incentives that lowers emissions, but does so in a manner that fails to align economic signals with

environmental goals. In allowing regulatory ambiguity to substitute for policy, governments forgo the dual advantages of targeted abatement and public revenue generation, ultimately sacrificing both efficiency and fiscal leverage.

2 Extraction and Exploration Equilibrium under Regulatory Uncertainty

Consider a panel of risk-neutral firms operating in the oil and gas sector, indexed by $k = 1, \dots, K$. Let $\mathcal{I} = \{1, 2, \dots, N\}$ denote the set of all oil&gas fields — with typical element i — and $\mathcal{I}^k \subseteq \mathcal{I}$ with $\mathcal{I}^k \cap \mathcal{I}^j = \emptyset$ for all $j \neq k$ be the subset of oil&gas fields in \mathcal{I} that are owned by firm k . Firms compete in quantities and are price-takers. The latter assumption is consistent with the negligible degree of market power asserted by Norwegian oil producers, and can be easily relaxed (see [Masnadi et al. \(2021\)](#)). The per-period field-level profit function is given by:

$$\Pi_t^{ik} = P_t^{ik} Q_t^{ik} - Cost(Q_t^{ik}, R_{t-1}^{ik}, S_t^{ik}, Reg_t, \varepsilon_t^{ik}) - W_t^{ik}, \quad (1)$$

where P_t^{ik} is the price at which firm k sells oil from field i , measured in real U.S. dollars per barrel (USD/BBL), and Q_t^{ik} is the quantity of oil extracted in that period, measured in million barrels per year (MMBBL/Year). The term R_{t-1}^{ik} denotes the available stock of recoverable reserves at the beginning of period t , expressed in million barrels (MMBBL). The variable S_t^{ik} is a dimensionless index capturing the stringency of the firm's internal environmental standards, while Reg_t is a similarly dimensionless index reflecting the stringency of national environmental regulation. The term W_t^{ik} represents firm-level exploration investment, measured in real U.S. dollars (USD), and ε_t^{ik} is a stochastic cost shock with zero mean and finite variance. Firms in period 0 do not observe the future realizations of Reg_t , and instead treat it as a random variable drawn from a known distribution, which may exhibit temporal dependence and need not be identically distributed over time. We further assume that this distribution is independent of firms' production and investment decisions, as well as of any private information available to firms at the time of the initial decision.

We assume that each firm k take all its decisions $\left\{ \left\{ Q_t^{ik}, W_t^{ik}, R_{t-1}^{ik}, L_{t-1}^{ik}, S_t^{ik} \right\}_{t=1}^{\infty} \right\}_{i \in \mathcal{I}^k}$

simultaneously at the beginning of period 1. Within this framework, firm k solves the following intertemporal optimization problem:

$$\begin{aligned} & \max_{\{Q_t^{ik}, S_t^{ik}, R_t^{ik}, L_t^{ik}\}_{t=1}^{\infty} \in X} \mathbb{E} \left[\sum_{t=1}^{\infty} \kappa^{t-1} \sum_{i \in \mathcal{I}^k} \left(P_t^{ik} Q_t^{ik} - Cost(Q_t^{ik}, R_{t-1}^{ik}, S_t^{ik}, Reg_t, \varepsilon_t^{ik}) - W_t^{ik} \right) \mid \Omega_0^k \right] \\ & \text{s.t. } R_t^{ik} = R_{t-1}^{ik} - Q_t^{ik} + D(W_{t-1}^{ik}, L_{t-1}^{ik}, \xi_t^{ik}) \\ & L_t^{ik} = L_{t-1}^{ik} + D(W_{t-1}^{ik}, L_{t-1}^{ik}, \xi_t^{ik}), \quad \text{for all } i \in \mathcal{I}^k, \text{ and } t = 1, 2, \dots \end{aligned}$$

where $\kappa \in (0, 1)$ is the discount factor, $\Omega_0^k = \{R_0^{ik}, L_0^{ik}, \{\xi_t^{ik}\}_{t=1}^{\infty}\}_{i \in \mathcal{I}^k}$ denotes firm k 's initial information set in period 0, $\mathbb{E}[\cdot]$ represents the expectation operator, and X is a compact set. The function $D(\cdot)$ is weakly increasing in W_{t-1}^{ik} , concave, and takes weakly positive values. It captures the quantity of new reserves discovered through exploration, which depends on lagged exploration investment W_{t-1}^{ik} , lagged cumulative exploration effort L_{t-1}^{ik} , and a stochastic innovation ξ_t^{ik} with finite mean and variance. Note that the information set Ω_0^k contains not only the initial conditions of the dynamic system but also the entire trajectory of future realizations of the random variable ξ_t^{ik} , meaning that each firm possesses perfect foresight regarding future shocks on discoveries⁴. The first constraint describes the law of motion for recoverable reserves at oilfield i : the stock in period t equals the previous period's reserves minus current extraction, plus newly discovered reserves. The second constraint governs the evolution of cumulative exploration effort, which increases with each period's discoveries — these, in turn, are a function of both current and past exploration intensity.

Our main theoretical contribution is to microfound the role of regulatory costs within the firm's cost function, allowing us to analyze how changes in *environmental regulatory risk* affect profit-maximizing behavior. To achieve this goal, we assume that the function $Cost(\cdot)$

⁴This assumption is imposed for technical tractability. Since all firm decisions are made in period $t = 0$, assuming perfect foresight of ξ_t^{ik} ensures that constraints are not violated with positive probability in future periods. This assumption does not influence the model's core trade-offs.

is strictly convex in $Q_t^{ik}, R_{t-1}^{ik}, S_t^{ik}$. This restriction ensures that the firm's optimization problem is well-behaved. Moreover, we assume that $Cost(\cdot)$ is additively separable in ε_t^{ik} . Since S_t^{ik} is chosen simultaneously with all other endogenous variables controlled by firm k , solely affects the per-period objective function in period t and does not enter any law of motion, and because firm's costs are additively separable across time and fields, it is possible to solve the firm's problem in two stages. Let \mathcal{S} be a compact subset of \mathbb{R} . The first stage consists in deriving a new conditional cost function $C(Q_t^{ik}, R_{t-1}^{ik}, \varepsilon_t^{ik})$, defined as the minimum over $S_t^{ik} \in \mathcal{S}$ of the expected cost, where the expectation is, in principle, conditional on all the endogenous choices made by firm k . However, thanks to the additive separability of firm's total costs across time periods and oilfields, as well as the additive separability of ε_t^{ik} in $Cost(\cdot)$, it is easy to show that the problem of interest is equivalent to one in which the aforementioned expectation is solely conditional on $\Omega_0^k, Q_t^{ik}, R_{t-1}^{ik}, \varepsilon_t^{ik}$. As a result, we can define $C(\cdot)$ as follows:

$$C(Q_t^{ik}, R_{t-1}^{ik}, \varepsilon_t^{ik}) = \min_{S_t^{ik} \in \mathcal{S}} \mathbb{E} \left[Cost(Q_t^{ik}, R_{t-1}^{ik}, S_t^{ik}, Reg_t, \varepsilon_t^{ik}) \mid \Omega_0^k, Q_t^{ik}, R_{t-1}^{ik}, \varepsilon_t^{ik} \right], \quad (2)$$

where the strict convexity in S_t^{ik} of $Cost(\cdot)$ and compactness of \mathcal{S} together ensure that a solution S_t^{ik*} to the problem above exists and is unique. Substituting S_t^{ik*} into the objective function of problem (2), we obtain the formula for $C(Q_t^{ik}, R_{t-1}^{ik}, \varepsilon_t^{ik})$, which can be used to solve the second stage problem.

2.1 Firm's Optimality Conditions

Incorporating the newly defined cost function $C(Q_t^{ik}, R_{t-1}^{ik}, \varepsilon_t^{ik})$, we can express firm k 's profit maximization problem as follows:

$$\begin{aligned} & \max_{\{Q_t^{ik}, R_t^{ik}, L_t^{ik}\}_{t=1}^{\infty} \in \mathcal{X}} \mathbb{E} \left[\sum_{t=1}^{\infty} \kappa^{t-1} \sum_{i \in \mathcal{I}^k} \left(P_t^{ik} Q_t^{ik} - C(Q_t^{ik}, R_{t-1}^{ik}, \varepsilon_t^{ik}) - W_t^{ik} \right) \mid \Omega_0^k \right] \\ \text{s.t. } & R_t^{ik} = R_{t-1}^{ik} - Q_t^{ik} + D(W_{t-1}^{ik}, L_{t-1}^{ik}, \xi_t^{ik}) \\ & L_t^{ik} = L_{t-1}^{ik} + D(W_{t-1}^{ik}, L_{t-1}^{ik}, \xi_t^{ik}), \quad \text{for all } i \in \mathcal{I}^k, \text{ and } t = 1, 2, \dots \end{aligned}$$

where \mathcal{X} is a compact set, and the corresponding Lagrangian writes:

$$\begin{aligned} \mathcal{L}_t^k = \mathbb{E} \left\{ \sum_{i \in \mathcal{I}^k} \sum_{t=1}^{\infty} \kappa^s [\Pi_t^{ik} + \lambda_t^{ik} [R_{t-1}^{ik} - R_t^{ik} - Q_t^{ik} + D_t^{ik}] + \right. \\ \left. + \mu_t^{ik} [L_t^{ik} - D_t^{ik} - L_{t-1}^{ik}]] \mid \Omega_0^k \right\}. \end{aligned}$$

It is possible to show that this second-stage optimization problem is well-behaved under the assumptions stated in the previous section (details in [Masnadi et al. \(2021\)](#)). Thus, we can solve for the first-order conditions to obtain the shadow prices of interest. The shadow price of discovered oil in field i in period t ,

$$\mathbb{E}[\lambda_t^{ik} \mid \Omega_0^k] = \mathbb{E} \left[P_t^{ik} \mid \Omega_0^k \right] - \mathbb{E} \left[\frac{\partial C_t^{ik}(\cdot)}{\partial Q_t^{ik}} \mid \Omega_0^k \right], \quad (3)$$

is obtained by solving the firm's optimality condition with respect to Q_t^{ik} and is equal to the difference between the price at which the field sells its output and its marginal production costs.⁵ The closer its magnitude is to zero the more the decision of the field management shifts from “*how much should the field produce?*” (intensive margin) to “*should the field*

⁵The formula can be easily extended to the case in which firms exert market power, see [Masnadi et al. \(2021\)](#) and [Benini et al. \(2023\)](#).

keep producing or cease business operations?" (extensive margin).

Similarly, the expected shadow-price of undiscovered oil in field i in period t , has formula:

$$\mathbb{E}[\mu_t^{ik} | \Omega_0^k] = \mathbb{E}[\lambda_t^{ik} | \Omega_0^k] - \left[\kappa \frac{\partial D_t^{ik}(\cdot)}{\partial W_t^{ik}} \right]^{-1}. \quad (4)$$

Equation (4) shows that the expected value of undiscovered oil equals the corresponding value of discovered oil minus the implicit cost of discovering an extra barrel in the following period, which is inversely proportional to the effectiveness of investment in discoveries captured by the derivative $\frac{\partial D_t^{ik}}{\partial W_t^{ik}}$. If its magnitude approaches zero, the decision of the field management shifts from “*how much should the firm invest in exploration for this field?*” (intensive margin) to “*should the field keep investing in exploration for this field or cease exploration operations?*” (extensive margin). Note that the assumption $\frac{\partial D_t^{ik}}{\partial W_t^{ik}} \geq 0$ implies that $\mathbb{E}[\mu_t^{ik} | \Omega_0^k] \leq \mathbb{E}[\lambda_t^{ik} | \Omega_0^k]$, meaning that firms typically cease to invest in exploration before they shut down business operation in a given field.

3 Empirical Analysis

3.1 Marginal Revenues

Business intelligence providers such as Rystad Energy offer detailed field-level data on production volumes and cost structures but do not report actual transaction prices between oil producers and refineries. As a result, we estimate expected field-level prices structurally. Building on the theoretical framework developed by [Benini et al. \(2023\)](#), we model the expected price as a function of a global benchmark — Brent, in the case of Norwegian oilfields — and the crude’s chemical characteristics, specifically its API gravity and sulfur content. The resulting pricing equation is:

$$\mathbb{E}[P_t^{ik} \mid \Omega_0^k] = \bar{P}_t + \beta_{1t}(API^{ik} - \overline{API}_t) + \beta_{2t}(S^{ik} - \bar{S}_t) , \quad (5)$$

where $\mathbb{E}[P_t^{ik} \mid \Omega_0^k]$ is the expected price at which field i sells its output in period t , measured in U.S. dollars per barrel (USD/BBL), conditional on the available information set Ω_0^k ; \bar{P}_t is the global benchmark price (Brent, in USD/BBL); API^{ik} and S^{ik} denote, respectively, the API gravity (unitless) and sulfur content (unitless, typically expressed in weight percent) of the crude produced by field i ; and \overline{API}_t and \bar{S}_t are the average API gravity and sulfur content of the global crude pool at time t , approximated by the average characteristics of crude oil imported into U.S. refineries. The coefficients β_{1t} and β_{2t} represent the marginal effects of a one-unit deviation in API gravity and sulfur content from the global average, and are both expressed in USD/BBL per unit of characteristic — that is, the unit of account for β_{1t} is USD/BBL per API degree, and for β_{2t} it is USD/BBL per percentage point of sulfur content. In this formulation, the expected field-level price equals the Brent price adjusted by a quality premium (or discount) based on how much lighter and sweeter the crude is relative to the market average. Both coefficient might evolve over time, reflecting shifts in refinery demand patterns, processing technologies, or both. Higher demand for light products (i.e jet fuel or gasoline) could increase how strong a field’s price responds to

$(API^{ik} - \overline{API}_t)$. On the other hand, if refineries improve so they can transform heavier oils into light fuels without added cost, the price response to $(API^{ik} - \overline{API}_t)$ may decrease. Lastly, a combination of these two effects may occur.

We cannot estimate equation (5) directly, as the dependent variable — field-level transaction prices — is not observed. However, our theoretical derivation of marginal revenues implies that the price of each internationally traded oil class equals the expected weighted average of the prices of the fields composing it. Following XXX, we assume that oilfields are classified in classes according to their API gravity and sulfur content. Let \mathcal{Z} denote a set of oilfields —possibly owned by different firms— such that $\{j, i\} \subseteq \mathcal{Z} \rightarrow (API^{ik} = API^{jl}) \cap (S^{ik} = S^{jl})$. That is, each class only includes oilfields producing oil with identical characteristics. Assuming that all private information is publicly available (i.e., $\Omega_0^k = \Omega_0^{\text{pub}}$), the price of oil class \mathcal{Z} at time t satisfies:

$$P_t(\mathcal{Z}) = \mathbb{E} \left[\sum_{j \in \mathcal{Z}} w_t^j P_t^{jk} \mid \Omega_0^{\text{pub}} \right] = \mathbb{E} \left[\sum_{j \in \mathcal{Z}} w_t^j P_t^{jk} \mid \Omega_0^k \right], \quad (6)$$

where \mathcal{Z} denotes a publicly traded oil class, $\{w_i^t\}_{i=1}^{N(z)}$ are time-varying weights reflecting the relative importance of each field i within class z , $N(z)$ the number of traded classes, and $P_t(z)$ is the observed spot price of oil class z at time t . Using data on twenty-two publicly traded oil classes — including spot prices, API gravity, and sulfur content — we construct a panel dataset, which spans from 1995 to 2022, and estimate the following regression:⁶

$$P_t(z) = \overline{P}_t + \beta_{1t}(API(z) - \overline{API}_t) + \beta_{2t}(S(z) - \overline{S}_t) + \nu_t, \quad (7)$$

where $\mathbb{E}[\nu_t \mid API(z), S(z)] = 0$. Before estimating equation (7), we assess the time-series properties of both the dependent variable $P_t(z)$ and the benchmark \overline{P}_t . Not all oil classes

⁶See Appendix A for a detailed description of the data sources, processing steps, and summary statistics.

have sufficiently long price histories to support formal testing, but for a subset of ten traded blends⁷, we conduct Augmented Dickey-Fuller (ADF) tests (Said and Dickey, 1984). In all cases, we fail to reject the null hypothesis of a unit root, indicating non-stationarity. The same result holds for the benchmark series \bar{P}_t . Given that both sides of equation (7) are integrated of order one, we transform the specification by subtracting \bar{P}_t from both sides. This yields to a specification, $P_t(z) - \bar{P}_t = \beta_{1t}(API(z) - \overline{API}_t) + \beta_{2t}(S(z) - \bar{S}_t) + \nu_t$, where the left-hand side is now a stationary series, consistent with standard results in time-series econometrics (Hamilton, 2020). This transformation allows us to estimate the structural parameters β_{1t} and β_{2t} using standard techniques.

Table 1 reports the average of the yearly conditional modes — i.e., the Best Linear Unbiased Predictors (BLUPs) — of the random coefficients (β_{1t}, β_{2t}) , estimated using the `ranef` function from the `lme4` package in R (Swamy, 1970; Bates et al., 2005; De Boeck et al., 2011; Bates et al., 2015). The full set of year-specific estimates is provided in Appendix ??.

Column (1) shows that in a specification where the deviation between the oil class price and the benchmark is explained solely by API gravity, a one-degree increase in API is associated, on average, with a price premium of 0.366 USD/BBL. Column (2) considers sulfur content as the sole regressor and finds that a one-percentage-point increase in sulfur reduces the price by 3.329 USD/BBL. In Column (3), where both characteristics are included, the estimated effect of API gravity drops to a more modest 0.009 USD/BBL, while the effect of sulfur remains virtually unchanged at -3.301 USD/BBL. All models are estimated using robust standard errors and report the Nakagawa-adjusted R^2 , which increases from 0.29 in Model 1 to 0.35 in Model 3 (Nakagawa and Schielzeth, 2013). We conclude that allowing for a time varying impact of the chemical characteristics of the crude explains 35% of the variance in the deviation between the price of an oil class and the benchmark.

⁷Cabinda, Bow River Heavy, Lloydminster, Oriente, Mayan, Bonny Light, Arabian Light, Arabian Medium, WTI, and Brent

Table 1. Price Regressions

	$P_t(z) - \bar{P}_t$ (USD/BBL)		
	Model 1	Model 2	Model 3
Δ API	0.37*** (0.30)		0.01 (0.01)
Δ S		-3.33*** (2.21)	-3.30*** (2.21)
Observations	345	345	345
R^2	0.29	0.35	0.35

Note: Nakagawa and Schielzeth (2013) standard errors in parentheses. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Under the assumptions laid out in this section, the estimators $\hat{\beta}_{1t}$ and $\hat{\beta}_{2t}$ are unbiased for the true structural parameters β_{1t} and β_{2t} . As such, they can be used to recover the unobserved expected transaction price at which individual Norwegian oilfields sell their output. For example, consider Ekofisk — one of Norway’s most significant oilfields in terms of cumulative production. Ekofisk crude has an API gravity of 39 and a sulfur content of 0.17%. In 2016, the Brent benchmark price was $\bar{P}_t = 48.95$ USD/BBL, while the average chemical characteristics of crude oil imported by U.S. refineries were $\overline{API}_t = 31.46$ and $\bar{S}_t = 1.4\%$. Using the estimated coefficients for that year, $\hat{\beta}_{1t} = 0.02$ and $\hat{\beta}_{2t} = -1.39$, the expected price at which Ekofisk sold its output is:

$$\begin{aligned}
\mathbb{E}[P_t^{\text{Ekofisk}} \mid \Omega_0^{\text{pub}}] &= \bar{P}_t + \hat{\beta}_{1t}(API^{\text{Ekofisk}} - \overline{API}_t) + \hat{\beta}_{2t}(S^{\text{Ekofisk}} - \bar{S}_t) \\
&= 48.95 + 0.02 \cdot (39 - 31.46) - 1.39 \cdot (0.0017 - 0.014) \\
&= 49.39 \text{ USD/BBL}
\end{aligned}$$

3.2 Marginal Costs

We assume that the firm’s cost function is additively separable between two components:

$$Cost(Q_t^{ik}, R_{t-1}^{ik}, S_t^{ik}, Reg_t, \varepsilon_t^{ik}) = PC(Q_t^{ik}, R_{t-1}^{ik}, \varepsilon_t^{ik}) + RC(Q_t^{ik}, S_t^{ik}, Reg_t), \quad (8)$$

where $PC(\cdot)$ represents standard production costs, which depend on the volume of oil extracted, the stock of recoverable reserves, and an idiosyncratic cost shock ε_t^{ik} ; whereas $RC(\cdot)$ denotes regulatory compliance costs, which depend on the extracted quantity, the stringency of the firm’s self-imposed environmental standards, and the level of national environmental regulation. Our analysis of production costs follows the standard functional form introduced by [Pesaran \(1990\)](#) and widely adopted in the literature ([Masnadi et al., 2021](#)):

$$PC(Q_t^{ik}, R_{t-1}^{ik}, \varepsilon_t^{ik}) = \zeta_0 + \zeta_1 Q_t^{ik} + \zeta_2 Q_t^{ik2} + \zeta_3 \frac{Q_t^{ik2}}{R_{t-1}^{ik}} + \varepsilon_t^{ik}. \quad (9)$$

The linear term Q_t^{ik} reflects standard variable costs that scale proportionally with output, such as labor, maintenance, and consumables. The quadratic term Q_t^{ik2} and the pressure-adjusted term $\frac{Q_t^{ik2}}{R_{t-1}^{ik}}$ jointly capture the nonlinear and reserve-dependent nature of marginal extraction costs. The former accounts for technical frictions that arise as production scales up — for example, congestion in surface facilities, diminishing returns to labor and equipment, or the need to mobilize additional infrastructure to sustain higher throughput. However, these frictions are further exacerbated by subsurface dynamics: as cumulative extraction depletes a reservoir, natural pressure declines, making it harder to sustain any given flow rate. The term $\frac{Q_t^{ik2}}{R_{t-1}^{ik}}$ captures this interaction by scaling the quadratic cost penalty inversely with the remaining stock of recoverable reserves. This formulation reflects the physical reality that a high production rate is more costly to achieve in a depleted field than in one with ample pressure support, due to increased reliance on energy-intensive techniques such as artificial lift, reinjection, or water flooding. Together, these terms allow the model to account for both surface and subsurface constraints on cost-efficiency.

In contrast to the production cost component, our main theoretical innovation lies in how we model regulatory compliance costs. To ensure empirical tractability, we introduce a functional form for regulatory costs that satisfies weak separability between firm-level and national-level regulatory stringency. Specifically, we assume that the pair $\{S_t^{ik}, Reg_t\}$ forms a weakly separable group (WSG) within the regulatory cost function $RC(\cdot)$; that is, the ratio

$$\frac{\partial RC(\cdot)}{\partial S_t^{ik}} \bigg/ \frac{\partial RC(\cdot)}{\partial Reg_t}$$

remains constant with respect to Q_t^{ik} over the relevant domain. This assumption implies that a firm's optimal self-regulation level S_t^{ik*} depends only on national regulatory conditions Reg_t , and not on production levels or other firm-specific variables. Based on this structure, we adopt the following functional form for $RC(\cdot)$:

$$RC(Q_t^{ik}, S_t^{ik}, Reg_t) = a(Q_t^{ik}) \cdot c(S_t^{ik}, Reg_t) , \quad (10)$$

where both $a(\cdot)$ and $c(\cdot)$ are strictly increasing and strictly convex C^2 functions. The multiplicative structure of the regulatory compliance cost represents a parsimonious way to ensure that the *WSG* assumption is satisfied, while simultaneously allowing for complementarities between operational scale and regulatory exposure. The latter feature captures the idea that the overall burden of regulation increases jointly with the scale of extraction and the stringency of regulation. That is, a higher volume of extracted oil may amplify the impact of stricter environmental standards — both firm-level and national — on total compliance costs.⁸ We derive parametric specifications for the functions $a(\cdot)$ and $c(\cdot)$ using second-order Taylor expansions. This approach yields the following functional forms:

$$a(Q_t^{ik}) = \delta_0 + \delta_1 Q_t^{ik} + \delta_2 Q_t^{ik2} , \quad (11)$$

⁸This desirable property would not hold under stronger assumptions like additive separability, making our formulation a practical balance between flexibility and tractability.

and

$$c(S_t^{ik}, Reg_t) = \alpha_0 + \alpha_1 S_t^{ik} + \alpha_2 Reg_t + \alpha_3 S_t^{ik2} + \alpha_4 Reg_t^2 + \alpha_5 S_t^{ik} Reg_t, \quad (12)$$

for appropriately chosen parameters $\alpha_0, \alpha_1, \dots, \alpha_5$, and $\delta_0, \delta_1, \delta_2$. By substituting the optimal self-regulation level,

$$S_t^{ik*} = -\frac{\alpha_1}{2\alpha_3} - \frac{\alpha_5}{2\alpha_3} \mathbb{E}[Reg_t | \Omega_0^k],$$

into the compliance cost function, we obtain an expression for regulatory costs:

$$RC(Q_t^{ik}, S_t^{ik*}, Reg_t) = \left[\delta_0 + \delta_1 Q_t^{ik} + \delta_2 Q_t^{ik2} \right] \cdot [\gamma_{0t} + \alpha_4 \mathbb{V}(Reg_t)], \quad (13)$$

where the first bracket captures the scale of production, and the second bracket isolates the influence of regulatory conditions. Specifically, γ_{0t} summarizes the effects of the first and second moment of the national regulation, while the term $\alpha_4 \mathbb{V}(Reg_t)$ captures the additional cost imposed by climate policy uncertainty.

3.2.1 Approximating $\mathbb{V}(Reg_t)$ via the Norwegian Climate Policy Uncertainty Index

To empirically operationalize the theoretical framework outlined above, we require a proxy for the variance of national environmental regulation, $\mathbb{V}(Reg_t)$. To this end, we construct a Climate Policy Uncertainty (CPU) index for Norway, following the methodology developed by [Berestycki et al. \(2022\)](#), which extends the original approach by [Baker et al. \(2016\)](#). This framework quantifies policy uncertainty by tracking the frequency of newspaper articles containing specific keyword combinations — typically “economic,” “policy,” and “uncertainty” — without relying on opaque machine learning techniques. Despite its simplicity, the method has proven robust, yielding results that closely align with major historical events [Baker et al. \(2016\)](#).

To adapt this framework to the context of climate policy, we develop an expanded lexicon of more than 40 climate-related keywords, capturing the broader vocabulary required for

topic-specific uncertainty measurement. Our keyword set builds on the OECD’s English-language CPU index, translated and refined to reflect the linguistic and policy nuances of the Norwegian context. Terms are grouped into three semantic categories — “climate,” “policy,” and “uncertainty” — and combined to identify relevant newspaper articles.⁹ Following the OECD’s methodology, we collect articles from three major national newspapers: *VG*, *Aftenposten*, and *Dagens Næringsliv*, accessed via the Retriever media archive.¹⁰ These outlets ensure broad coverage across general, political, and economic reporting. While following this procedure, we simultaneously adapt the keyword list to the national context, minimize false positives, and mitigate the risk of false negatives¹¹. First, adapting international vocabularies to national media requires careful linguistic and contextual judgment. For example, the English keyword “hydro” translates broadly as “vann” in Norwegian, which is too vague for our purposes. We therefore opted for more specific terms such as “vannkraft” (hydropower) and “vannmagasin” (hydro reservoir). Second, false positives — irrelevant articles containing valid keywords — can distort the index. To limit these, we manually reviewed hundreds of randomly sampled articles and iteratively refined our search algorithm. Words that frequently triggered unrelated matches were excluded. For instance, “miljø” (environment) often appears in unrelated terms like “arbeidsmiljø” (work environment), and “politisk klima” (political climate) refers to political discourse rather than climate change. Third, false negatives — relevant articles not captured by the search — are harder to assess directly. However, as noted by [Berestycki et al. \(2022\)](#), this is unlikely to bias the index unless the share of undetected articles varies systematically with macroeconomic conditions. We follow the OECD in assuming that such variation is minimal.

To make the resulting index comparable across time and newspapers, we follow the standardization procedure used by both [Baker et al. \(2016\)](#) and [Berestycki et al. \(2022\)](#).

⁹A detailed list of search terms, along with quarterly and monthly versions of the index, is provided in Appendix B.

¹⁰See Appendix B.2 for a detailed discussion of newspaper selection and archive choice. Appendix B.3 discusses the limitations of using Retriever.

¹¹See Appendix B.4 for a more detailed discussion on the challenges of false positives and false negatives

The goal of this approach is to capture meaningful within-country variation in climate policy uncertainty, while minimizing noise introduced by structural differences across media sources. We begin by retrieving, for each of the three newspapers, the annual number of articles flagged by our CPU search algorithm, as well as the total number of articles published that year. This allows us to compute a raw frequency ratio for each outlet:

$$\text{Ratio}_t^j = \frac{\text{CPU Articles}_t^j}{\text{Total Articles}_t^j},$$

where j indexes the newspaper and t the year. These raw ratios reflect the share of CPU-related content in each outlet’s coverage. However, as emphasized by Baker et al. (2016), raw ratios are not directly comparable across newspapers or over time due to differences in editorial practices, article volume, and content structure. To address this, we standardize each newspaper-specific time series by dividing the raw ratio by its own standard deviation over the full sample period. We then average the standardized series across the three newspapers to construct a single, unified time series. Finally, we normalize this aggregated series to have a mean of 100 over the full time period,

$$CPU_t = \left(\frac{1}{3} \sum_{j=1}^3 \frac{\text{Ratio}_t^j}{\text{SD}(\text{Ratio}^j)} \right) \cdot \left(\frac{100}{\text{Mean} \left(\frac{1}{3} \sum_{j=1}^3 \frac{\text{Ratio}_t^j}{\text{SD}(\text{Ratio}^j)} \right)} \right). \quad (14)$$

Figure 1 shows that Norwegian climate policy uncertainty remained relatively stable until the mid-2000s, before rising steadily from around 2011. Volatility in the index increases sharply after the 2015 Paris Agreement, reflecting a period of intensified debate and ambiguity around Norway’s climate transition. While there is broad political consensus on the importance of meeting climate targets, there is less agreement on how to achieve them. Key episodes contributing to spikes in uncertainty include the political shift following the 2013 election and subsequent reinterpretation of the 2012 Climate Agreement; contentious debates over biofuel policy and electric vehicle subsidies in 2017–2018; and disagreements

over wind power, particularly regarding whether profitable onshore wind should still receive subsidies and whether unprofitable offshore wind should be publicly supported. A major peak in 2019 corresponds to intense political debate over whether to open ecologically sensitive areas near Lofoten, Vesterålen, and Senja to oil and gas exploration — a proposal that was ultimately rejected, marking a significant victory for environmental advocates, local communities, and the Labour Party. The final major spike occurs during the 2021–2022 energy crisis, when surging electricity prices reignited controversy over the electrification of offshore oil platforms and its implications for domestic energy demand.




Figure 1. Authors’ own calculation, based on the A-tekst archives with data from 1995 to 2023. The search algorithm to generate this index can be found in Appendix [A](#)

The observed volatility in the index — particularly in the years following the Paris Agree-

ment — suggests that firms increasingly face not just evolving regulations, but uncertainty about how these regulations will unfold. To formally incorporate this notion into our model, we interpret the CPU index as a proxy for the variance of future regulatory outcomes:

$$CPU_t \cong \chi_0 + \chi_1 \mathbb{V}(Reg_t) + \chi_{3t} ,$$

where χ_{3t} captures transitory or unobserved variation. This assumption reflects the design of the index as a measure of regulatory *risk* rather than regulatory *intensity* — that is, the second rather than the first moment of the policy distribution.

Substituting the formulas for $PC(\cdot)$, $a(\cdot)$, $c(\cdot)$, CPU_t and S_t^{ik*} into (8) and making use of the previously stated restrictions on the distribution of Reg_t , we obtain the empirical specification of the cost function, which writes:

$$C_t^{ik} = \theta_{0t} + \theta_{1t} Q_t^{ik} + \theta_{2t} Q_t^{ik2} + \theta_{3t} \frac{Q_t^{ik2}}{R_{t-1}^{ik}} + \theta_{4t} CPU_t + \theta_{5t} CPU_t Q_t^{ik} + \varepsilon_t^{ik} . \quad (15)$$

To estimate the empirical cost function in equation (15), we compile a field-level panel dataset that includes operational expenditures, annual production, recoverable reserves, and climate policy uncertainty¹². These data are sourced from Rystad Energy’s Ucube Upstream database, which provides detailed information on both producing and decommissioned fields. We begin by excluding any field with zero cumulative production, as these do not reflect active extraction behavior. We further restrict the sample to fields with nonzero output in at least one year between 1995 and 2023, since our model requires observed production to estimate cost parameters. In addition, we focus only on years in which production is strictly positive, under the assumption that climate policy uncertainty does not meaningfully affect cost decisions at inactive sites. After applying these filters, we obtain an unbalanced panel covering 128 unique fields and 1,658 field-year observations over the period 1995–2023.

¹²All monetary values, including OPEX from Rystad, are deflated using the same CPI index employed in the estimation of marginal revenues.

Table 2. Summary Statistics for Cost Variables

Variable	Unit of Account	Mean	Median	SD	Min	Max
C	Million USD/Year	181.46	99.70	240.12	0.00	1776.09
Q	Million BBL/Year	14.01	5.00	23.90	0.00	193.59
R	Million BBL	189.57	52.25	341.71	0.03	2383.61
CPU	Pure Number	135.92	114.75	96.90	27.49	347.53

Note: Data are based on 128 unique oilfields observed over the period 1995–2023.

The summary statistics in Table 2 reveal substantial heterogeneity in the cost and production characteristics of oilfields on the Norwegian Continental Shelf. Annual operational expenditures (OPEX) average 181.5 million USD, but the median is significantly lower at 99.7 million USD, reflecting a highly skewed distribution driven by a subset of high-cost fields. The maximum value exceeds 1.77 billion USD, while a large number of fields operate with considerably lower expenditure. This dispersion is consistent with the presence of technical complexity and varying field maturity across the shelf. Production volumes (Q) also display wide variation. The mean annual output per field is 14.01 million BBL, but the median is only 5.00 million BBL, again highlighting the dominance of a few large producers. Fields such as Troll, Gullfaks, and Oseberg — particularly active in the late 1990s and early 2000s — substantially raise the average. In contrast, fields like Sindre and Enoch, developed after 2009, contribute only marginally to aggregate output, with the latter having initial reserves as low as 2.29 million BBL compared to over 2.64 billion BBL for Oseberg. Recoverable reserves (R) follow a similarly skewed distribution. The average field holds about 189.6 million BBL in reserves, but the median is closer to 52.3 million BBL. The standard deviation is 341.7 million BBL, and the range extends from just 0.03 million BBL to more than 2.38 billion BBL, illustrating stark differences in geological endowments and field scale. This distribution reflects the typical discovery pattern in petroleum geology: many small

fields interspersed with a few “super-giant” basins. Finally, the Climate Policy Uncertainty (CPU) index captures sharp changes in the regulatory environment over time. With a mean of 135.9, a median of 114.8, and a standard deviation of 96.9, the index ranges from 27.5 to 347.5 across the sample. As visualized in Figure 1, CPU levels remained relatively stable until the early 2010s, after which both the level and volatility increased markedly. The sharpest spikes occur after the 2015 Paris Agreement, coinciding with political debates over climate commitments, the 2018 Climate Settlement, and controversies surrounding exploration in ecologically sensitive areas such as Lofoten, Vesterålen, and Senja.

Like in the case of the pricing equation, we estimate the cost function using a linear mixed-effects model with time-varying coefficients $(\theta_{0t}, \theta_{1t})$. Table 3 reports the average of the yearly conditional modes — i.e., the Best Linear Unbiased Predictors (BLUPs) — of the random coefficients, computed using the `ranef` function from the `lme4` package in R. The full set of year-specific estimates is available in Appendix H.

Table 3. Cost Regressions

	ΔC (Million USD/Year)		
	Model 1	Model 2	Model 3
ΔQ	15.68*** (0.80)	15.74*** (0.79)	15.22*** (0.72)
ΔQ^2	-0.01* (0.00)	-0.01* (0.00)	-0.01** (0.00)
$\Delta(Q^2/R)$	16.03*** (3.31)	16.02*** (3.31)	14.48*** (3.32)
ΔCPU		-0.00 (0.09)	-0.05 (0.09)
$\Delta(\text{CPU} \times Q)$			0.01** (0.00)
Observations	1,581	1,581	1,581
R^2	0.15	0.15	0.16

Note: Nakagawa Schielzeth (2013) standard errors in parentheses. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Across all three specifications, we observe consistent and statistically significant patterns. The linear term in production, ΔQ , is positive and highly significant in every model, with point estimates ranging from approximately 15.2 to 15.7 million USD per additional million barrels produced. This reflects the strong cost elasticity with respect to output — a result consistent with the expectation that variable operational costs rise with production. The quadratic term ΔQ^2 is negative and weakly significant (10% level in Models 1 and 2; 5% in Model 3), suggesting mild diminishing returns to scale at high production levels.

However, the magnitude of this effect is modest. In contrast, the term $\Delta(Q^2/R)$, which captures pressure-related cost dynamics, is positive and significant at the 1% level across all specifications. This confirms the theoretical prediction that producing large volumes from depleted fields is more expensive due to the need for artificial lift or enhanced recovery technologies as reservoir pressure declines. The CPU enters the cost function in Models 2 and 3. While the standalone effect ΔCPU is small and statistically insignificant, suggesting that the level of uncertainty has no effect on the level of the cost for a given level of environmental regulation, the interaction term $\Delta(\text{CPU} \times Q)$ in Model 3 is positive and significant at the 5% level. This indicates that the marginal cost of production increases more steeply in periods of greater policy uncertainty — consistent with the idea that regulatory risk compounds the operational challenges of scaling up output. In other words, high production levels are more costly in uncertain policy environments. Overall model fit remains moderate, with R^2 values around 0.15–0.16, which is typical for field-level cost regressions. These results align well with the theoretical structure of the cost function and provide robust empirical support for key mechanisms — including production effects, reservoir depletion, and uncertainty-driven cost amplification.

3.3 Marginal Discoveries

Lastly, our analysis of exploration activity and the discovery of new oil reserves follows standard practice in the literature [Masnadi et al. \(2021\)](#). We assume that the firm’s discovery function takes the additive form:

$$D_t^{ik} = \gamma_0 + \gamma_1 W_{t-1}^{ik} + \gamma_2 L_{t-1}^{ik} + \gamma_3 L_{t-1}^{ik,2} + \xi_t^{ik} , \quad (16)$$

where D_t^{ik} denotes the volume of oil discovered at field i in year t (measured in million barrels per year, MM BBL/year), W_{t-1}^{ik} is exploration investment in the previous year (in

million USD)¹³, and L_{t-1}^{ik} is the cumulative amount of oil previously discovered at field i up to year $t-1$ (in MM BBL). The error term ξ_t^{ik} captures unobserved drivers of exploration success. Equation (16) satisfies the theoretical conditions under the restrictions $\gamma_1 \geq 0$, $\gamma_2 \leq 0$, and $\gamma_3 \leq 0$, which reflect intuitive economic relationships. The condition $\gamma_1 \geq 0$ ensures that higher investment increases the likelihood of new discoveries. The terms $\gamma_2 \leq 0$ and $\gamma_3 \leq 0$ imply diminishing returns to exploration in mature oilfields: the more oil that has already been discovered at a site, the less likely it is that additional exploration will yield large new finds. Formally, this is captured by the marginal effect of cumulative discoveries,

$$\frac{\partial D_t^{ik}}{\partial L_{t-1}^{ik}} = \gamma_2 + 2\gamma_3 L_{t-1}^{ik} \leq 0, \quad \text{and} \quad \frac{\partial^2 D_t^{ik}}{\partial^2 L_{t-1}^{ik}} = 2\gamma_3 \leq 0 .$$

We restrict our analysis to oilfield, which have positive exploration capital expenditures ($W_{t-1}^i > 0$), and past cumulative discoveries ($L_{t-1}^i > 0$). Given the aging profile of the Norwegian Continental Shelf, the number of new discoveries in recent years is limited. The final dataset comprises 347 field-year observations across 88 unique fields, with descriptive statistics for the discovery function variables reported in Table 4.

Table 4. Summary Statistics for Discoveries Variables

	Unit of Account	Mean	Median	SD	Min	Max
D	Million BBL / Year	1.57	0.77	1.79	0.00	7.49
W	Million USD / Year	2.41	0.00	15.66	0.00	147.85
L	Million BBL	35.17	23.19	41.14	0.14	206.49

Note: Data are based on 88 unique oilfields over the period 1995–2023.

Exploration capital expenditures are highly skewed across fields. The median field records virtually no spending in a given year, while the mean is 2.41 million USD and the maximum

¹³All monetary values, including exploration CAPEX from Rystad, are deflated using the same CPI index employed in the estimation of marginal revenues and costs.

exceeds 147 million USD. This pattern reflects the episodic nature of exploration activity, where firms concentrate investment in select years when geological or market conditions are favorable, leading to a small number of high-expenditure episodes that inflate the average. A similar skew is evident in the distribution of annual discoveries. While the median field adds approximately 0.77 million BBL of recoverable reserves per year, the mean rises to 1.57 million BBL, driven by occasional large discoveries — such as the 2022 addition of nearly 7.2 million BBL at the Martin Linge field. Most fields experience modest yearly additions, punctuated by infrequent but substantial reserve finds. Cumulative discoveries also exhibit significant heterogeneity: values range from just 0.14 million BBL in newer or marginal fields to over 206 million BBL in legacy fields like Gullfaks. The median field has accumulated roughly 23.2 million BBL, while the mean is considerably higher at 35.2 million BBL, underscoring the presence of a few dominant fields. This distributional asymmetry is consistent with well-documented patterns in petroleum geology, where the discovery landscape is shaped by many small finds and a few giant ones.

Unlike the pricing equation and the cost function, equation (16) does not feature time-varying coefficients. We estimate it using Ordinary Least Squares (OLS) on the first-differenced form of the model, implemented via the base `lm` function in R. We apply a first-difference transformation because the dependent variable (discovered oil volumes) is strictly positive and exhibits a skewed distribution, violating the normality assumption. Differencing mitigates this issue by stabilizing the variance and reducing skewness, thereby avoiding the need for a generalized linear model with non-constant marginal effects. See Table 5 for results.

Table 5. Discovery Regressions

	Δ Discoveries (Million BBL/Year)		
	Model 1	Model 2	Model 3
ΔW	0.02 (0.01)	0.02 (0.01)	0.02 (0.01)
ΔL		-0.30** (0.13)	-0.04 (0.13)
ΔL^2			-0.004** (0.00)
Observations	347	347	347
R^2	0.01	0.01	0.01

Note: Robust standard errors in parentheses. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Table 5 reports the results of three first-differenced OLS specifications of the discovery function. Across all models, the coefficient on changes in exploration investment (ΔW) is positive, stable, and statistically insignificant at conventional levels. This suggests a weak empirical link between year-to-year variation in exploration expenditures and new reserve additions, possibly due to the lumpy nature of exploration activity or reporting noise. In Model 2, the coefficient on cumulative discoveries (ΔL) is significantly negative, supporting the hypothesis of diminishing geological potential as fields mature. However, this relationship becomes nonlinear in Model 3, where the quadratic term (ΔL^2) is also negative and statistically significant at the 5% level. This confirms decreasing returns to exploration: as fields become more geologically exhausted, the likelihood of meaningful new discoveries falls at an increasing rate. Despite these theoretically consistent signs, the explanatory power of all models remains modest, with R^2 values around 0.01, indicating

that much of the variance in discovery outcomes remains unexplained. This is likely due to unobserved geological heterogeneity and the episodic nature of exploration success, which is only partially captured by field-level inputs. Overall, the results provide suggestive — though not definitive — evidence that both investment intensity and field maturity shape exploration outcomes in economically and geologically plausible ways.

4 Welfare Analysis

4.1 Shadow Price of Discovered Oil

Microeconomic Analysis In Table 6, we compare the estimated shadow price of discovered oil, defined as

$$\mathbb{E}[\hat{\lambda}_t^{ik} | \Omega_0^k] = \left[\bar{P}_t + \hat{\beta}_{1t}(API^i - \overline{API}_t) + \hat{\beta}_{2t}(S^i - \bar{S}_t) \right] - \left[\hat{\theta}_{1t} + 2\hat{\theta}_2 Q_t^{ik} + 2\hat{\theta}_3 \frac{Q_t^{ik}}{R_{t-1}^{ik}} + \hat{\theta}_5 CPU_t \right],$$

with its theoretical counterpart obtained by replacing CPU_t with its minimum observed value, $\min(CPU_t)$. Under the minimum climate policy uncertainty assumption, the mean shadow price of discovered oil across our sample is approximately 69.09 USD/BBL, with a median of 63.34 USD/BBL and a standard deviation of 37.56 USD/BBL. By contrast, when we allow CPU_t to vary across its observed range, the mean estimated shadow price decreases slightly to 68.39 USD/BBL, with a median of 61.39 USD/BBL and an almost unchanged standard deviation. The minimum and maximum values remain broadly consistent. The difference between the two pricing regimes has a mean of 0.70 USD/BBL and a median of 0.38 USD/BBL, implying that climate policy uncertainty reduces the marginal value of discovered oil by approximately 1.01%.

Table 6. Summary Statistics for the Shadow Price of Discovered Oil

	Unit of Account	Mean	Median	SD	Min	Max
$\mathbb{E}[\hat{\lambda}_t^{ik, \min} \Omega_0^k]$	USD/BBL	69.09	63.34	37.56	0.53	142.01
$\mathbb{E}[\hat{\lambda}_t^{ik} \Omega_0^k]$	USD/BBL	68.39	61.39	37.55	0.28	141.64
$\Delta\lambda$	USD/BBL	0.70	0.38	0.70	0.00	2.24

Note: Data are based on 128 unique oilfields observed over the period 1995–2022.

Although this effect is economically meaningful, its magnitude increases when we restrict

the sample to post-2015 observations¹⁴ — i.e., the period surrounding the Paris Agreement and the adoption of new emissions targets, the EU Green Deal, and related regulatory developments. In this restricted period, the mean shadow price under minimum CPU falls to 57.11 USD/BBL, with a median of 57.41 USD/BBL, and a notably lower standard deviation of 17.38 USD/BBL. This reduction reflects, among other things, the lower average oil prices and changes in cost structures following the 2014 oil price collapse and subsequent market adjustments. When CPU is allowed to vary within the post-2015 sample, the mean shadow price declines further to 55.74 USD/BBL, with a median of 55.47 USD/BBL. The implied reduction due to CPU averages 1.36 USD/BBL, with a median of 1.29 USD/BBL. This corresponds to a 2.38% reduction relative to the average value under minimum policy uncertainty. Compared to the full-sample average reduction of approximately 1.01%, this suggests that the marginal impact of climate policy uncertainty on investment decisions has intensified in recent years — potentially reflecting heightened regulatory scrutiny, more ambitious decarbonization targets, or increased investor sensitivity to long-term policy risks.

In addition to estimating shadow prices, we enrich our analysis by incorporating upstream carbon intensity data for a subset of 56 oilfields on the Norwegian Continental Shelf.¹⁵ These data are sourced from the Oil Production Greenhouse Gas Emissions Estimator (OPGEE), a bottom-up life cycle assessment model developed to estimate greenhouse gas emissions across the oil extraction process. Specifically, OPGEE quantifies upstream emissions — from exploration through to the refinery gate — by integrating field-specific geological, technological, and logistical parameters. By accounting for variation in reservoir characteristics, extraction methods, energy use, and associated gas management (e.g., flaring

¹⁴See tables with summary statistics for the 2015-2022 period in Appendix G

¹⁵The 56 fields for which we obtain carbon intensity data are broadly representative of the Norwegian upstream sector. Between 1995 and 2023, they account for an average of approximately 70% of total annual oil production on the Norwegian Continental Shelf, with yearly coverage ranging from a minimum of 47% to a maximum of 80%. Coverage ratios are calculated by aggregating annual production data from Rystad Energy’s UCube database and computing the share of total NCS output attributable to the matched fields.

and venting), OPGEE captures substantial heterogeneity in emissions intensity across fields. Combining OPGEE-based carbon intensity estimates with our shadow price data allows us to link each field’s carbon footprint to its estimated change in marginal value under climate policy uncertainty. This integrated dataset enables a novel environmental welfare analysis, in which we quantify the shadow price of carbon dividing the uncertainty costs by the measured upstream emissions. The results of our welfare analysis are presented in Tables 7–G.2.

The estimated difference in shadow prices for this subset of fields is slightly lower than that observed in the full sample. Specifically, the mean of the delta declines from 0.70 USD/BBL (with a median of 0.38 and a standard deviation of 0.70) to 0.55 USD/BBL (with a median of 0.31 and a standard deviation of 0.62). By multiplying this difference by annual production volumes, we quantify the fiscal cost associated with climate policy uncertainty. In essence, we estimate the revenue the Norwegian government could have captured had uncertainty remained at its minimum observed level and been replaced with an equivalent tax, holding the overall stringency of climate regulation constant. As reported in Table 7, the estimated average annual revenue shortfall amounts to approximately 6.5 million USD per oilfield. In the most stable years, this loss is negligible, while in years of heightened uncertainty it can rise to nearly 17.5 million USD per field. Then, we multiply the field-level (upstream) carbon intensity, measured in kg CO₂/BBL, by the annual production volumes to obtain an aggregate measure of emissions for each field-year observation. The resulting distribution of yearly emissions ranges from approximately 0.25 million tonnes of CO₂ to 2.39 million tonnes, with a mean of 0.87 million tonnes and a median of 0.61 million tonnes per field per year. To put these figures in perspective, a typical passenger vehicle emits approximately 4.7 tonnes of CO₂ per year.¹⁶ Accordingly, the average field-year in our sample emits as much CO₂ as roughly 185,000 cars (i.e., 0.87×10^6 tonnes \div 4.7 tonnes

¹⁶See U.S. Environmental Protection Agency, “Greenhouse Gas Emissions from a Typical Passenger Vehicle”, EPA-420-F-14-040a, May 2014. Available at: <https://nepis.epa.gov/Exe/ZyPURL.cgi?Dockey=P100JPPH.TXT>

per car). At the upper end of the distribution, the most carbon-intensive field-year reaches emissions equivalent to more than 500,000 cars, highlighting the substantial environmental impact of field-level upstream oil production.

These field-specific emissions estimates allow us to compute an implied shadow carbon tax by dividing the estimated uncertainty cost by the corresponding emissions level:

$$\text{Shadow Carbon Tax}_t^{ik} = \frac{\text{Uncertainty Cost}_t^{ik}}{\text{Carbon Emissions}_t^{ik}}.$$

As shown in our estimates, the implied shadow carbon tax averages approximately 15.4 USD per tonne of CO₂, with a median of 8.2 USD, a standard deviation of 17.7 USD, and a maximum of nearly 63 USD. These are substantial figures—particularly given that they do not arise from any explicit carbon pricing mechanism, but rather emerge endogenously from the fiscal penalties associated with climate policy uncertainty at the field level. The effect becomes even more pronounced in the post-2015 period¹⁷, following the Paris Agreement. During these years, the average shadow carbon tax more than doubles to 38.4 USD per tonne, while the minimum increases from zero to nearly 20 USD. This increase occurs despite a notable decline in average field-level emissions, reflecting reduced production. It is not emissions growth but sustained policy ambiguity that drives the escalation in implicit carbon costs. For comparison, the price of carbon allowances in the EU Emissions Trading System (ETS) has recently stabilized in the range of 70–90 EUR per tonne (approximately 75–95 USD). While the estimated average shadow tax remains below this benchmark, its magnitude is nonetheless remarkable: even in the absence of direct regulation, uncertainty alone imposes a de facto cost on emissions that can approach compliance market levels. This finding underscores the fiscal and behavioral significance of policy credibility. The shadow carbon tax offers a transparent, field-level metric for capturing the hidden economic cost of regulatory ambiguity. In this sense, uncertainty functions not merely as a risk factor, but as a shadow pricing mechanism with real effects on valuation, incentives, and environmental

¹⁷See Appendix G for full summary table

outcomes.

Table 7. Summary Statistics for the Welfare Effect

	Unit of Account	Mean	Median	SD	Min	Max
Uncertainty Cost	Million USD / Year	6.53	5.83	4.63	0.00	17.48
Carbon Emissions	Million Ton CO ₂ / Year	0.87	0.61	0.66	0.25	2.39
Shadow Carbon Tax	USD / Ton CO ₂	15.42	8.16	17.67	0.00	62.91

Note: Data are based on 56 unique oilfields observed over the period 1995–2022.

Macroeconomic Analysis To complement the field-level (microeconomic) analysis, we now aggregate uncertainty-related fiscal losses and upstream carbon emissions across all 56 fields for which carbon intensity data are available. This allows us to assess the broader economic and environmental implications of climate policy uncertainty. Over the 1995–2023 period, we estimate that foregone government revenues attributable to climate policy uncertainty average approximately 255 million USD per year, with a median of 215 million USD and a maximum exceeding 800 million USD in some years. The standard deviation — 212 million USD — reflects substantial inter-annual volatility, driven not only by field-level heterogeneity but also by shifts in market cycles, see Table 8.

Table 8. Summary Statistics for the Aggregate Welfare Effect

	Unit of Account	Mean	Median	SD	Min	Max
Fiscal Loss	Million USD / Year	254.72	215.11	212.37	0.00	803.84
Carbon Emissions	Million Ton CO ₂ / Year	31.64	29.71	16.42	12.30	58.78

Note: Data are based on the 1995–2022 annual time series.

To contextualize these losses, consider that Norway’s central government collects between 2,000 and 2,400 billion NOK per year in total revenue (roughly 200–240 billion

USD)¹⁸. On this scale, the annual fiscal loss from climate policy uncertainty corresponds to approximately 0.1–0.2% of total government revenue. While this represents a small share of the total budget, the average annual fiscal loss from climate policy uncertainty — 255 million USD — is equivalent to approximately 17% of the revenue collected from Norway’s net wealth tax, which generates around 15 billion NOK (roughly 1.5 billion USD) per year¹⁹. In high-uncertainty years, the estimated fiscal loss can exceed 50% of total wealth tax revenues. Viewed through this lens, the fiscal impact of climate policy uncertainty is far from negligible. Though it does not appear prominently in budgetary ledgers, it is quantitatively comparable to one of Norway’s most visible and contested sources of public revenue. This highlights the potential fiscal gains from strengthening regulatory clarity and policy credibility — without the distributional controversies typically associated with raising statutory tax rates.

¹⁸See Statistics Norway (SSB), *General government revenue and expenditure*, available at: <https://www.ssb.no/en/statbank/table/12082>

¹⁹See OECD (2022), *The Role and Design of Net Wealth Taxes in the OECD*, and Statistics Norway. See also Civita (2023), “Hvor mye får staten inn på formuesskatten?”




Figure 2. This shows how the yearly total fiscal penalty of climate policy uncertainty moves compared to the climate policy uncertainty index.

The total upstream carbon emissions across our sample average 31.6 million tonnes of CO₂ per year, with a median of 29.7 Mt and a standard deviation of 16.4 Mt. Annual values range from 12.3 Mt — during years of depressed activity — to nearly 59 Mt in expansionary phases, illustrating the strong linkage between emissions and production cycles, dampened by steady output from mature legacy fields and the management of co-extracted natural gas. To situate this magnitude, Norway’s total annual greenhouse gas emissions (across all sectors) were approximately 64 million tonnes of CO₂-equivalents in 2022, according to

Statistics Norway²⁰. Thus, our sample—focused solely on upstream petroleum—accounts for about 50% of national emissions. In environmental terms, this means that fluctuations in upstream production and climate-policy uncertainty within these 56 fields can swing Norway’s total emissions by over 30 Mt annually. Such a large share highlights the outsized impact of upstream oil activity on national CO₂ outcomes and underscores how policy uncertainty in this segment resonates powerfully at the macro level.

4.2 Shadow Price of Undiscovered Oil

Microeconomic Analysis Using the discovery function in Section 2, we also estimate the shadow price for fields that have yet to be drilled. By using the structural coefficients from our discovery regressions, we estimate the discovery function presented in Section 2, $\mathbb{E}[\mu_t^{ik} | \Omega_0^k] = \mathbb{E}[\lambda_t^{ik} | \Omega_0^k] - \left[\kappa \frac{\partial D_t^{ik}(\cdot)}{\partial W_t^{ik}} \right]^{-1}$. Our dataset covers the 1995-2022 period and 50 fields, as some fields drop out when we rely on available data for discoveries to estimate $\left[\frac{\partial D_t^{ik}(\cdot)}{\partial W_t^{ik}} \right]^{-1}$.

Starting with Table 9, when climate policy uncertainty is fixed at its historical minimum, the field-level shadow prices of undiscovered oil μ_{min}^{CPU} has a mean of 25.46 USD/BBL, with a median of 14.89 USD/BBL and a substantial standard deviation of 38.94 USD per BBL. Value ranges from a very low -44.27 USD/BBL up to 91.82 USD/BBL for the max, thus indicating that there are some fields that are clearly not worth developing, while others are highly attractive prospects. Allowing for full time varying climate policy uncertainty, μ^{CPU} is slightly reduced to a mean of 24.71 USD/BBL, and a median of 12.92 USD/BBL, while the SD is virtually unchanged. Furthermore, the minimum to maximum range is fairly similar to the same range for μ_{min}^{CPU} , and is estimated at -44.52 USD/BBL up to 91.45 USD/BBL. The increased effect of climate policy uncertainty is estimated to average about 0.75/per BBL, indicating that the marginal value discovering one more barrel of oil is

²⁰Statistics Norway (2025), “Emissions from Norwegian economic activity,” Table 09298: 63,297 kt CO₂-eq in 2023; see also total greenhouse gas emissions series 13932.

reduced by an average of 0.75 USD/BBL taking climate policy into account when valuating potential prospects. Half the fields experience a reduction of 0.64 USD per BBL, while the max reduction reaches up to 2.24 USD per BBL. Although modest in the mean, the wide underlying distribution shows that several exploration targets remain highly sensitive to regulatory debate - much as we observed for shadow prices of discovered oil.

Table 9. Field-level Shadow Price of Undiscovered Oil

	Unit of Account	Mean	Median	SD	Min	Max
μ_{\min}	USD/BBL	25.457	14.892	38.942	-44.272	91.818
μ	USD/BBL	24.708	12.916	38.993	-44.520	91.452
$\Delta\mu$	USD/BBL	0.748	0.639	0.654	0.000	2.240

Note: Summary statistics. $\Delta\mu^{CPU} = \mu_{\min}^{CPU} - \mu^{CPU}$. Data are based on 50 fields on the Norwegian Continental Shelf.

Again, to complement our field-level analysis for undiscovered oil, we run similar welfare calculations as for our discovered oil resources. Results are reported in Table 10. We make a small change when calculating the fiscal cost, as we multiply $\Delta\mu$ by discoveries instead of production. For the estimate of emissions we multiply the carbon intensity of each field by its discoveries. This follows from the fact that the volumes of interest when estimating the shadow price of undiscovered oil are discoveries, due to the fact that there generally is no production in the field yet. Table 10 report an estimated average annual fiscal loss of only 0.8 million USD per oil field. That figure is an order of magnitude smaller than the 6.5 million USD reported for producing fields, reflecting the modest scale and sporadic nature of new discoveries. Even so, the standard deviation of approximately 0.73 million USD and the maximum average fiscal loss stands at 3.32 million USD per oilfield. This underscores that in peak exploration years, regulatory noise linked to climate policy uncertainty still translates into non-trivial fiscal losses. Because discoveries add reserves

rather than barrels at the wellhead, the associated emissions are minuscule: mean annual emissions attributable to undiscovered resources are just 0.07 Mt CO_2 , representing about 3 % of production related emissions (Table 7). Despite trivial fiscal and emissions exposure, the implicit carbon tax is already of the same order of magnitude as for operating assets at a mean of approximately 17 USD per tonne of CO_2 - indicating that climate policy uncertainty pushes hurdle rates long before a field moves into development and production. The average shadow carbon tax ranges from 0 to more than 77 USD per tonne, and this maximum value is closely aligned with the current price of carbon allowances in the EU Emissions Trading System.

To gauge how tightening global climate ambition has reshaped exploration incentives, we replicate the calculations for the Paris-era subsample. Results, drawn from Table G.4, reveal a striking escalation with an increase of more than 80% for the cost of uncertainty estimated at almost 1.5 million USD per year. In this period we also note a 45% reduction in carbon emissions, which also likely stems from a significantly lower rate of new discoveries in the last decade. Still, this leads to a spike in the estimated shadow carbon tax, which doubles to 38.8 USD per tonne CO_2 .

Table 10. Yearly Summary Statistics, Average Across Fields

	Unit of Account	Mean	Median	SD	Min	Max
Uncertainty Cost	Million USD / Year	0.799	0.780	0.731	0.000	3.315
Carbon Emissions	Million Ton CO_2 / Year	0.070	0.055	0.044	0.012	0.192
Shadow Carbon Tax	USD / Ton CO_2	16.786	10.039	19.048	0.000	77.447

Note: Data are based on 50 unique oil fields observed over the period 1995–2022.

Macroeconomic Analysis. Applying the same aggregation used for producing fields, Table 11 shows that uncertainty-related fiscal losses for exploration prospects average USD 9 million per year (median \approx 3.9 million), spiking to 49.7 million USD when discovery

activity coincides with intense policy debate. That is still barely 4 % of the 255 million USD hit sustained by producing assets, and—with a much tighter dispersion ($SD \approx 1.2$ million versus 212 million)—it highlights how sporadic finds create only modest fluctuations in the government’s exposure. Annual CO₂ emissions attributable to new discoveries average 0.50 Mt, or roughly 2 % of the 31.6 Mt emitted by operating fields; their narrow range (0.06–0.90 Mt, $SD \approx 0.27$ Mt) confirms that exploration adds a minor, low-volatility slice to the basin’s carbon ledger. In short, the government currently faces limited fiscal and environmental risk from deferred exploration, yet lingering policy ambiguity is already pushing up hurdle rates for tomorrow’s barrels and, by extension, shaping the long-run trajectory of Norwegian Continental Shelf development.

Table 11. Yearly Summary Statistics, Total Sum Across Fields

	Unit of Account	Mean	Median	SD	Min	Max
Fiscal Loss	Million USD / Year	9.007	3.865	1.197	0.000	49.732
Emission Reduction	Million Ton CO ₂ / Year	0.495	0.361	0.265	0.062	0.904

Note: Data are based on the 1995–2022 time series.




Figure 3. This shows how the yearly total fiscal penalty of climate policy uncertainty moves compared to the climate policy uncertainty index.

Climate policy uncertainty bites harder in the post-Paris Agreement period even for barrels that have yet to be found, though the absolute stakes remain modest. Table G.4 shows the average CPU-driven wedge in the shadow price of undiscovered oil jumps from about 0.6 USD / BBL pre-Paris to 1.3 USD / BBL post-Paris (range between 0.67–2.24 USD). This mirrors the increase observed for producing assets, $\Delta\lambda$ in Table 7 and Table G.1, indicating that heightened policy debate after 2015 is equally suppressing exploration incentives. Average field-level Fiscal losses due to uncertainty cost rises but stays small in this period. Because discovery volumes are limited, uncertainty cost climbs to only 1.5 million USD per year on average—dwarfed by the 11 million USD post-Paris loss for

producing fields (Table G.1). Even so, the upper tail reaches USD 3.3 million in active exploration years²¹. Furthermore, annual emissions attributable to discoveries fall to 0.04 Mt CO_2 , roughly half the long-run average and just a quarter of one percent of post-Paris production-related emissions (15 Mt). Table G.5 aggregates the post-Paris welfare impact of climate-policy uncertainty for undiscovered resources and confirms that, while the absolute numbers remain tiny next to producing assets, the stakes have increased materially since the 1995-2022 average reported in Table 11.

Annual fiscal losses now averages about 20 million USD (median is approximately 19.6 million USD), more than double the long-run undiscovered figure of 9 million USD and roughly 4 percent of the 521 million USD shortfall associated with producing fields in the same 2015-2022 window (Table G.2). The maximum reaches 49.7 million USD, reflecting years in which exploration activity and regulatory turbulence coincide; the standard deviation (13.6 million USD versus 184 million USD for producing oilfields) underscores the still sporadic nature of discoveries on the Norwegian Continental Shelf. Looking at emissions we see that total CO_2 attributable to new discoveries averages 0.50 Mt per year, essentially unchanged from the 0.49 Mt long-run mean and a mere 3 % of the 15 Mt emitted by producing fields post-Paris. With a narrow range of 0.21–0.90 Mt and a standard deviation of approximately 0.27 Mt, exploration-related emissions remain a minor, low-volatility component of the basin’s carbon intensity.

Although undiscovered resources still represent only a minor fraction of the basin-wide fiscal and environmental exposure, post-Paris era climate policy uncertainty is already eroding into the economics of these resources. Annual uncertainty related revenue losses have risen from roughly 9 million USD to 20 million USD, and the per-barrel CPU penalty matches that of producing oilfields: decision-makers would willingly accept a premium of 1-2 USD/BBL for a clear and stable policy environment. Individually modest, these premiums accumulate across high-risk prospects and can redirect exploration capital. The

²¹See Table G.4 in Appendix G

results echo our finding for producing oilfields: credible, predictable climate policy is rapidly becoming a prerequisite not just for safeguarding existing output, but for justifying the search for tomorrow's resources.

5 Conclusion

This paper develops a structural framework to study how climate policy uncertainty (CPU) affects the valuation of oil resources and the efficiency of extraction and exploration decisions. Using a novel CPU index constructed from Norwegian media sources, we show that even in a politically stable country like Norway, regulatory ambiguity has intensified since 2011, with pronounced spikes after the 2015 Paris Agreement. We integrate this measure into a field-level model of oil production and discovery on the Norwegian Continental Shelf, allowing CPU to influence regulatory compliance costs and, consequently, the marginal cost of extraction. Our analysis yields field-specific shadow prices for both discovered and undiscovered oil, which serve as sufficient statistics for investment and production incentives. We find that higher CPU systematically depresses these values, with the strongest effects occurring post-2015.

We interpret this distortion as an implicit “shadow carbon tax,” ranging from \$15 to \$38 per tonne of CO₂. However, unlike a well-designed Pigouvian tax, this implicit price does not scale with emissions intensity and fails to generate fiscal revenue. It penalizes producers indiscriminately and inefficiently, distorting investment and reducing emissions without rewarding low-emitting technologies or financing green transition policies. The result is an uncoordinated and opaque mechanism that sacrifices both economic efficiency and public finance. Our findings suggest that clarifying the trajectory of climate policy could simultaneously reduce economic deadweight loss, enhance fiscal outcomes, and support a more predictable path toward decarbonization.

References

- Baker, S. R., Bloom, N. and Davis, S. J. (2016), ‘Measuring economic policy uncertainty’, *The quarterly journal of economics* **131**(4), 1593–1636.
- Basaglia, P., Carattini, S., Dechezleprêtre, A. and Kruse, T. (2021), ‘Climate policy uncertainty and firms’ and investors’ behavior’, *Unpublished manuscript* .
- Bates, D., Mächler, M., Bolker, B. and Walker, S. (2015), ‘Fitting linear mixed-effects models using lme4’, *Journal of statistical software* **67**, 1–48.
- Bates, D. et al. (2005), ‘Fitting linear mixed models in r’, *R news* **5**(1), 27–30.
- Benini, G., Brandt, A. R., Dotti, V. and El-Houjeiri, H. (2023), ‘The economic and environmental consequences of the petroleum industry extensive margin’, *University Ca’ Foscari of Venice, Dept. of Economics Research Paper Series No. 14* .
- Berestycki, C., Carattini, S., Dechezleprêtre, A. and Kruse, T. (2022), ‘Measuring and assessing the effects of climate policy uncertainty’.
- Dagbladet (2024), ‘- Det verste scenarioet for Norge — dagbladet.no’, <https://www.dagbladet.no/nyheter/det-verste-scenarioet-for-norge/81783518>. [Accessed 14-08-2024].
- De Boeck, P., Bakker, M., Zwitser, R., Nivard, M., Hofman, A., Tuerlinckx, F. and Partchev, I. (2011), ‘The estimation of item response models with the lmer function from the lme4 package in r’, *Journal of Statistical Software* **39**, 1–28.
- EIA (2024), ‘Eia short-term energy outlook’, <https://www.eia.gov/dnav/pet/hist/LeafHandler.ashx?n=PET&s=I000000008&f=M>. [Accessed 19-03-2024].
- EIA (n.d.), ‘Eia short-term energy outlook’, <https://www.eia.gov/outlooks/steo/realprices/>, year = 2024, note = [Accessed 19-03-2024],.
- Gavriilidis, K. (2021), ‘Measuring climate policy uncertainty’, *Available at SSRN 3847388* .
- Hamilton, J. D. (2020), *Time series analysis*, Princeton university press.
- Ifwarsson, E. J. V., Kleiven, L. E., Sendstad, L. H. and Hagspiel, V. (2021), ‘Policy uncertainty in scandinavian countries’, *Economics Letters* **209**, 110121.
- Masnadi, M. S., Benini, G., El-Houjeiri, H. M., Milivinti, A., Anderson, J. E., Wallington, T. J., De Kleine, R., Dotti, V., Jochem, P. and Brandt, A. R. (2021), ‘Carbon implications of marginal oils from market-derived demand shocks’, *Nature* **599**(7883), 80–84.

- Nakagawa, S. and Schielzeth, H. (2013), ‘A general and simple method for obtaining r^2 from generalized linear mixed-effects models’, *Methods in ecology and evolution* **4**(2), 133–142.
- Nordhaus, W. D. (2007), ‘A review of the stern review on the economics of climate change’, *Journal of economic literature* **45**(3), 686–702.
- Pesaran, M. H. (1990), ‘An econometric analysis of exploration and extraction of oil in the uk continental shelf’, *The Economic Journal* **100**(401), 367–390.
- Said, S. E. and Dickey, D. A. (1984), ‘Testing for unit roots in autoregressive-moving average models of unknown order’, *Biometrika* **71**(3), 599–607.
- Shiller, R. J. (2020), *Narrative economics: How stories go viral and drive major economic events*, Princeton University Press.
- Stern, N. (2006), ‘Stern review: The economics of climate change’.
- Swamy, P. A. (1970), ‘Efficient inference in a random coefficient regression model’, *Econometrica: Journal of the Econometric Society* pp. 311–323.
- Weitzman, M. L. (2007), ‘A review of the stern review on the economics of climate change’, *Journal of economic literature* **45**(3), 703–724.

Appendices

A Appendix A: Expected Price Estimation

A.1 Data Sources for Price Regressions

We make use of the U.S. Energy Information Agency’s (EIA (2024)) open data to access the average price of landed crude in the United States. In addition we also use the open source data from EIA (EIA (n.d.)) to obtain the Consumer Price Index which we use to adjust our nominal prices to real prices. For our data we use 2023 as the base year for the CPI adjustment. We choose to use Brent Crude as our global benchmark price as Brent is the most widely used benchmark price in Norway, we download the Brent prices for Rystad Energy’s Ucube before we adjust them for CPI to turn the prices from nominal to real. For our estimate of \overline{API}_t and \overline{S}_t we make use of **double check with Giacomo where this data comes from.**, while our API^i and S^i , corresponding to the Norwegian oil fields in our data sample, are sourced from Rystad Energy’s Ucube.

A.2 Summary Statistics for Expected Prices

Table A.1 presents summary statistics for API gravity and sulphur content of Norwegian crude oil fields and U.S. average API and sulphur values, using data spanning 1995–2023. The Norwegian crude oil data reflects considerable variability between fields, with API gravity averaging 38.507 and a median closely aligned at 38.500, along with a relatively high standard deviation of 7.058, indicating diverse oil qualities across fields on the Norwegian Continental Shelf. The sulphur content for Norwegian oil is comparatively low, indicating a sweet average for Norwegian oil. Sulphur levels are averaging 0.325% , yet it displays substantial variation from as low as 0.02% to as high as 3.57%. In contrast, the U.S. data represent national average values, thus varying primarily over time rather than across fields. The U.S. average API gravity is lower and exhibits less variability, with a mean of 31.110

and a small standard deviation (0.847). Similarly, the U.S. average sulphur content is higher, averaging 1.361% and showing limited temporal variation, indicated by the small standard deviation of 0.094. These distinct patterns underscore differences in oil characteristics between Norwegian fields and average U.S. crude oil over time, significantly influencing their respective market prices and refining processes.

Table A.1. Summary Statistics for API and Sulphur

	Unit of Account	Mean	Median	SD	Min	Max
$API(z)$	Pure Number	38.507	38.500	7.058	19.000	58.000
$S(z)$	Percentage	0.325	0.230	0.356	0.020	3.570
\overline{API}_t	Pure Number	31.110	30.985	0.847	30.180	33.210
\overline{S}_t	Percentage	1.361	1.405	0.094	1.130	1.460

Note: API is measured in degrees API gravity; Sulphur in weight percent. Data cover the period 1995–2023. Norwegian field-level API ($API(z)$) and sulphur ($S(z)$) are from Rystad Energy; U.S. averages \overline{API}_t and \overline{S}_t serve as benchmarks.

Table A.2 summarizes two key price variables used in our empirical model. The benchmark price, \overline{P} —proxied by Brent and entering equation (x)—averages 78.32 USD/BBL over 1995–2023, with a median of 72.04 USD/BBL a standard deviation of 37.63, and a range spanning 23.93 USD to 150.72 USD. By contrast, the expected field-level selling price, P_t^{ik} , exhibits a slightly higher mean of 81.88 USD/BBL and greater dispersion (SD of 39.18 USD; range between 19.21 USD/BBL and 159.48 USD/BBL), reflecting quality premia specific to individual fields on the Norwegian Continental Shelf. The broader spread and higher mean of P_t^{ik} relative to \overline{P}_t underscore how local market conditions can deviate from global benchmarks, while the pronounced right tail in both series captures price spikes during crisis periods, such as per example the 2008 commodity boom and the 2022 energy crisis. The gap between mean and median in each case also indicates moderate positive

skewness, consistent with occasional but pronounced surges in crude prices over the sample period.

Table A.2. Summary Statistics for Price Variables

	Unit of Account	Mean	Median	SD	Min	Max
\bar{P}_t	USD/BBL	78.321	72.036	37.628	23.926	150.722
P_t^{ik}	USD/BBL	81.876	70.589	39.182	19.205	159.475

Note: Sample covers 1995–2023. \bar{P}_t is the benchmark (Brent), and P_t^{ik} is the field-specific selling price.

A.3 Types Of Crude Oil

Crude oil is categorized into different types primarily based on two key properties: API gravity, which measures density, and sulphur content. Generally, oil types with higher API values (lighter oils) and lower sulphur content (sweeter oils) command higher market prices. Table A.3 provides a comprehensive overview of 22 distinct crude oil varieties sourced globally, spanning regions such as North and South America, the Middle East, West and Central Africa. Notably, Brent Crude from the United Kingdom represents the sole European crude type in our dataset.

The table highlights considerable price variation across these oil types, ranging from approximately 37.5 USD/BBL for the heavy, sulphur-rich Venezuelan Furrial crude, to as high as 110.5 USD/BBL for the light and sweet Saharan Blend from Algeria. Furthermore, significant price variability exists even within individual countries. Venezuela, for instance, produces multiple crude varieties: alongside Furrial, it also includes Leona and Merey, priced notably higher at averages of approximately 42.25 USD/BBL and 95.22 USD/BBL, respectively. These substantial variations underscore the diverse characteristics and corresponding market valuations of global crude oil types.

Table A.3. Summary Statistics: Prices for Oil Types

Oil Type	Country of Origin	Mean	Max	Min	SD	API	Sulphur
Arabian Light	Saudi Arabia	71.73	148.24	23.10	36.00	32.80	1.97
Arabian Medium	Saudi Arabia	69.39	145.11	20.30	37.11	30.20	2.59
Basrah Light	Iraq	100.26	144.50	50.66	35.32	30.50	2.90
Berri	Saudi Arabia	103.71	150.06	57.92	35.35	38.50	1.50
Bonny Light	Nigeria	76.44	158.43	25.46	39.85	33.40	0.16
Bow River Heavy	Canada	56.79	119.32	19.46	27.76	24.70	2.10
Brent	United Kingdom	75.46	150.72	23.93	36.56	38.30	0.37
Forcados Blend	Nigeria	65.17	157.22	26.82	35.09	30.80	0.16
Furrial	Venezuela	37.43	54.24	22.88	8.58	30.00	1.06
Leona	Venezuela	42.25	112.59	18.30	21.31	24.00	1.50
Light Sour Blend	Canada	91.03	130.75	50.83	28.73	64.00	3.00
Lloydminster	Canada	56.66	116.78	18.97	28.90	20.90	3.50
Marlim	Brazil	109.17	151.72	61.42	36.61	19.60	0.67
Mayan	Mexico	63.01	133.86	17.22	34.34	21.80	3.33
Merey	Venezuela	95.22	137.07	49.47	34.09	15.00	2.70
Napo	Ecuador	93.07	134.75	47.56	34.84	19.00	2.00
Olmecca	Mexico	54.70	143.17	25.38	30.15	37.30	0.84
Oriente	Ecuador	70.52	140.02	21.59	35.48	24.10	1.51
Qua Iboe	Nigeria	135.03	158.53	96.93	26.56	36.30	0.14
Rabi-Kouanga	Gabon	56.64	135.13	25.51	30.46	37.70	0.15
Saharan Blend	Algeria	110.51	156.90	63.25	37.93	45.00	0.09
WTI	United States	74.09	140.40	26.90	32.49	39.60	0.24

Note: All prices are CPI-adjusted to 2023 U.S. dollars per barrel.

Figure A.1 provides a scatter plot illustrating the relationship between API gravity, sulphur content, and real crude oil prices per barrel (adjusted to 2016 dollars). API gravity is represented along the x-axis, while the CPI-adjusted oil price per barrel is displayed on the y-axis. Sulphur content is depicted through color gradients ranging from blue (lowest sulphur, or sweetest crude) to red (highest sulphur, or most sour crude). The visualization distinctly highlights how crude oil types with higher API gravity and lower sulphur content tend to command higher market prices. Specifically, widely recognized oil benchmarks such as Brent and West Texas Intermediate (WTI) are situated prominently among the highest-priced varieties, attributable to their favorable combination of approximately 40-degree API gravity and notably low sulphur content.

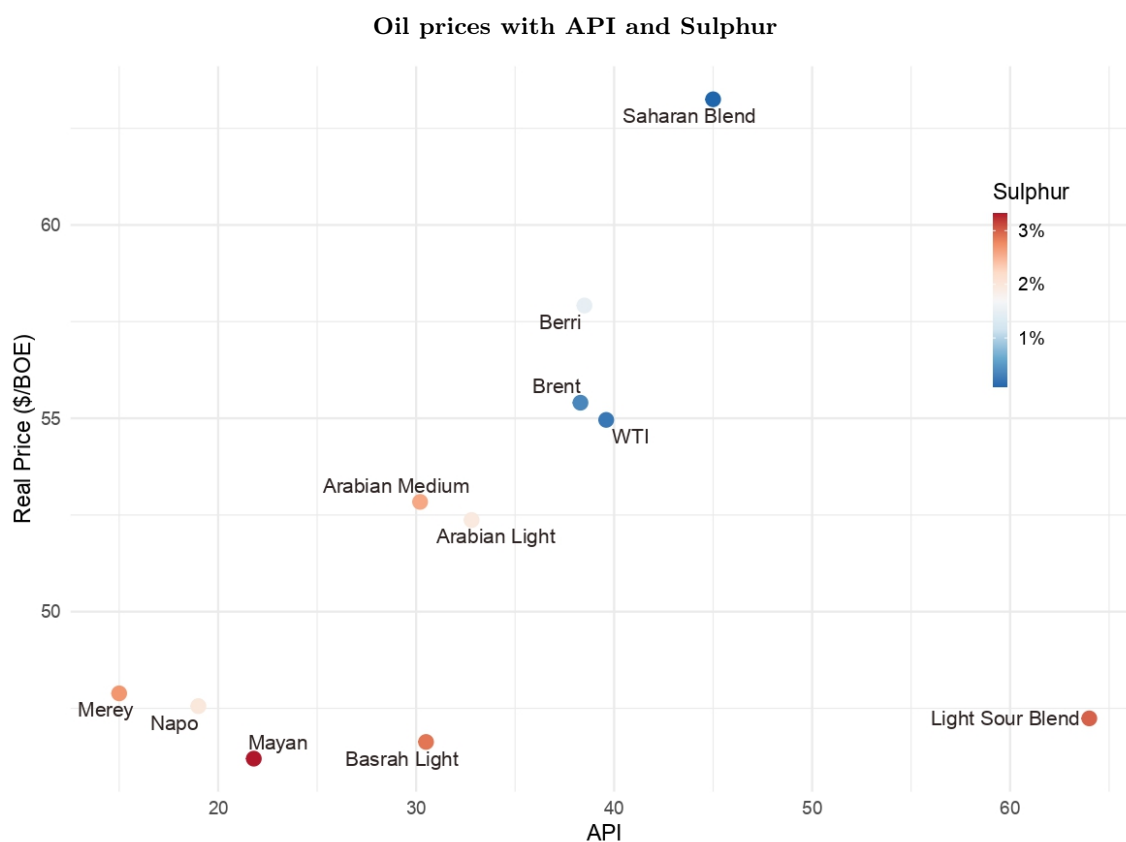


Figure A.1. This scatter plot is based on the year 2016

A.4 Oil Prices For Different Types Of Oil

Figure A.2 depicts the real-term evolution of oil prices from 1978 to 2023, adjusted to 2023 dollars using CPI data from the US Energy Information Administration (?). The graph clearly illustrates significant fluctuations driven by global supply and demand dynamics, geopolitical events, and economic conditions. Notably, prices exhibit pronounced volatility around 2008, coinciding with the global financial crisis and subsequent recovery, and again around 2014–2015, correlating with OPEC’s strategic shift to maintain market share, causing a substantial decline. Moreover, distinct price patterns among different crude oil

types become evident, with lighter and sweeter oils (such as Brent and WTI) consistently commanding higher prices relative to heavier, sourer crude varieties like Lloydminster and Mayan. This persistent price difference highlights how quality attributes, such as API gravity and sulphur content, critically influence oil valuation and market positioning over time.

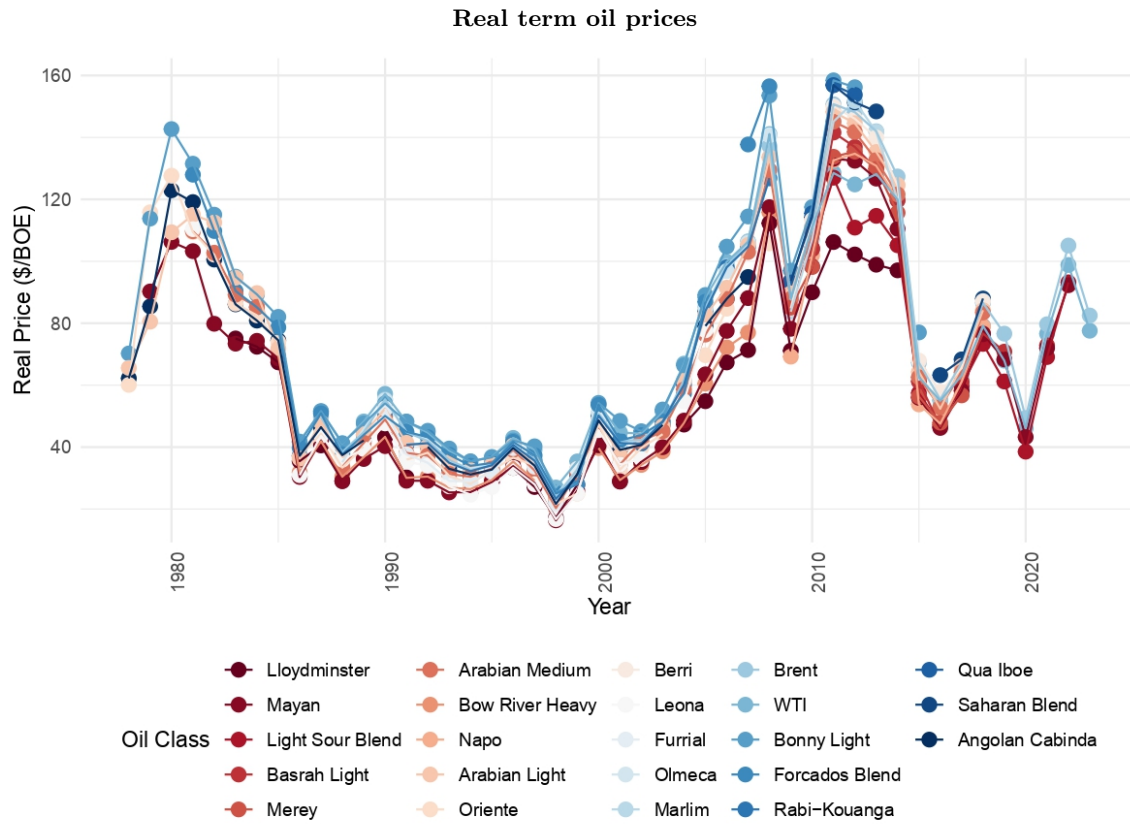


Figure A.2. The real prices are CPI adjusted for the CPI for 2023 based on the US Energy Information Administration (EIA) Short-term outlook

A.5 Norwegian API and Sulfur Confidence Intervals

In Table [A.4](#) display the values, standard deviation and confidence intervals for our ΔAPI and $\Delta Sulphur$ coefficients in Table [1](#).

Table A.4. Confidence Intervals for Δ API and Δ Sulphur

Year	Δ API				Δ Sulphur			
	Value	St. Dev.	CI (2.5%)	CI (97.5%)	Value	St. Dev.	CI (2.5%)	CI (97.5%)
1995	0.004	0.057	-0.108	0.116	-1.409	1.555	-4.457	1.638
1996	0.005	0.057	-0.107	0.117	-1.633	1.552	-4.676	1.409
1997	0.005	0.057	-0.108	0.117	-2.318	1.557	-5.370	0.733
1998	0.005	0.057	-0.108	0.117	-2.041	1.577	-5.131	1.050
1999	0.005	0.057	-0.107	0.117	-1.073	1.584	-4.177	2.032
2000	0.004	0.057	-0.108	0.116	-2.900	1.586	-6.007	0.208
2001	0.005	0.057	-0.107	0.117	-4.172	1.589	-7.287	-1.057
2002	0.005	0.057	-0.107	0.117	-2.055	1.592	-5.176	1.065
2003	0.003	0.057	-0.109	0.116	-2.768	1.592	-5.889	0.352
2004	0.006	0.057	-0.106	0.119	-3.901	1.605	-7.046	-0.756
2005	0.014	0.057	-0.099	0.126	-6.771	1.592	-9.891	-3.651
2006	0.014	0.057	-0.099	0.126	-7.065	1.592	-10.186	-3.944
2007	0.009	0.057	-0.103	0.121	-8.244	1.592	-11.365	-5.123
2008	0.007	0.057	-0.106	0.119	-7.437	1.592	-10.558	-4.317
2009	0.027	0.054	-0.079	0.133	-2.805	1.371	-5.493	-0.117
2010	0.032	0.054	-0.073	0.138	-3.422	1.360	-6.088	-0.756
2011	0.027	0.054	-0.079	0.132	-5.892	1.356	-8.549	-3.235
2012	-0.010	0.054	-0.116	0.096	-6.025	1.358	-8.687	-3.363
2013	0.009	0.054	-0.097	0.114	-4.937	1.404	-7.689	-2.185
2014	0.004	0.054	-0.103	0.110	-0.982	1.514	-3.950	1.986
2015	0.019	0.054	-0.087	0.126	-2.077	1.467	-4.952	0.797
2016	0.024	0.054	-0.083	0.130	-1.388	1.612	-4.547	1.771
2017	0.018	0.054	-0.088	0.124	-1.268	1.455	-4.119	1.583
2018	0.013	0.054	-0.093	0.120	-0.154	1.617	-3.323	3.016
2019	0.002	0.058	-0.111	0.115	-1.968	3.171	-8.184	4.247
2020	0.002	0.058	-0.111	0.115	-2.033	3.324	-8.548	4.481
2021	0.002	0.058	-0.111	0.115	-2.106	3.253	-8.483	4.270
2022	0.004	0.058	-0.109	0.117	-3.584	3.352	-10.154	2.985

Note: 95% confidence intervals for annual changes in API and sulphur.

B Appendix B: Search strategy for the Norwegian Climate Policy Uncertainty Index

We construct the Climate Policy Uncertainty (CPU) index for Norway following the methodology developed by [Berestycki et al. \(2022\)](#), which extends the approach originally proposed by [Baker et al. \(2016\)](#). Their framework measures economic policy uncertainty by analyzing newspaper articles containing a core set of keywords — typically the triplet “economic”, “policy”, and “uncertainty” — without relying on hard-to-disentangle black box machine learning algorithms. Despite its simplicity, this method has proven robust and aligns well with major historical events.²²

Our lexicon of search terms are divided into three categories; climate, policy, and uncertainty, in Norwegian it translates to; klima, politikk, usikkerhet. In the first category we have terms such as "klima", "karbon", "klimaendringer", "klimagassutslipp", which are linked to the overall climate, but we also include more specific industrial and technological terms such as "vannkraft", "vindmøller", "elbil". In the second category we use words linked to regulation and policy making such as "politikk", "regulering", "lovgivning", and also more specific terms related to design and development of climate policies such as "klimakvoter", "kvotemarked", "miljømerking", "grønne sertifikat". Finally, in the third category, we use words linked to uncertainty, such as "usikker", "usikkerhet", "uklar" and "uklart".

While [Baker et al. \(2016\)](#) use a relatively narrow keyword set, constructing a topic-specific index like CPU requires a more elaborate vocabulary. For English-speaking countries, the OECD includes over 60 climate-related terms in its CPU construction, compared to just 10 in the original index ([Berestycki et al., 2022](#)). We adopt this expanded keyword approach to build a CPU index tailored to the Norwegian context. To construct the index, we use

²²See www.policyuncertainty.com for a comprehensive overview.

over 40 climate-related search terms to balance breadth and precision — capturing as many relevant articles as possible while minimizing irrelevant noise. Our lexicon builds on the OECD’s Climate Policy Uncertainty vocabulary²³, which we translated and adapted to Norwegian, refining it to reflect the local language and context.

B.1 Final Search Strategy

((energi OR klima OR klimaendring* OR karbon OR "global oppvarming" OR klimagassutslipp OR utslipp* OR GHG OR karbondioksid OR CO2 OR metan OR CH4 OR forurensing OR vannkraft OR vannmagasin OR vindkraft OR vindmøller OR vindmølle OR vindpark OR solkraft OR elbil) NEAR/25(politikk* OR regulering OR reguleringer OR lovgivning OR lovgivninger OR direktiv OR direktiver OR avgift OR avgifter OR skatt OR skatter OR subsidie OR subsidier OR klimakvoter OR kvoter)) AND (usikker OR usikkerhet or uklar or uklart)

B.2 Newspaper and archive choices

We rely on the Retriever media archive,²⁴ adopting the OECD’s practice of basing uncertainty measures on widely read national outlets with broad readership. Our baseline index therefore draws on three high-circulation and reputable Norwegian newspapers—*VG*, *Aftenposten*, and *Dagens Næringsliv*. In 2022 these titles published 308 469, 252 448, and 94 094 articles, respectively, placing them at the top of general and business-news circulation rankings in Norway.²⁵ Two of these papers—*VG* and *Aftenposten*—were also used by [Ifwarsson et al. \(2021\)](#) in constructing Norway’s Economic Policy Uncertainty index, strengthening comparability across studies.

We replace their third outlet, *Finansavisen*, with *Dagens Næringsliv* because Retriever’s

²³See [Berestycki et al. \(2022\)](#), Annex A for the full list of English search terms.

²⁴Retriever is the largest media archive in the Nordics, covering newspaper archives from 1945 to the present.

²⁵According to Mediebedriftenes Landsforening, the Norwegian media association, *VG* and *Aftenposten* were the most widely circulated papers in 2022 with 308 469 and 252 448 articles, respectively; *Dagens Næringsliv* had 94 094 articles, the highest among business newspapers.

Atekst database offers far deeper historical coverage of the latter (from 1997 onward, versus only from 2011 for *Finansavisen*).²⁶ Robustness checks that expand the search to as many as ten newspapers yield highly correlated results with substantial article overlap;²⁷ hence the three-paper strategy provides a parsimonious yet comprehensive foundation for our Climate Policy Uncertainty index.

B.3 Limitations of Retriever / A-tekst

We face some challenges linked to using the Norwegian archive Retriever compared to Factiva, especially in regards to one of the restrictions. The OECD can impose a restriction that "climate" and "policy" words have to be in the same paragraph. This is, as far as we know, not yet possible in Retriever. To investigate whether this is a big challenge when constructing our index, we compared the OECD CPU index with the CPU index proposed by Gavriilidis (2021). Even though they use slightly different search strategies, like the fact that Gavriilidis (2021) does not impose a restriction that "climate" and "policy" terms must be in the same paragraph. We find that the two different CPU indices tend to correlate. The OECD index might provide slightly more precision with the stricter search restrictions, but as Retriever cannot apply all the similar restrictions, we find that excluding this restriction is a fair trade-off in the case for construction of the Norwegian Climate Policy Uncertainty index. Moreover, the US based CPU and UK based CPU²⁸ can rely on more comprehensive amounts of news articles compared to the Norwegian version. Thus, the Norwegian CPU might not be as precise, but it still works well as a solid proxy for Norwegian climate policy uncertainty as it still shows peaks around periods of heightened uncertainty regarding climate policies.

²⁶Based on Retriever's Atekst archive as of January 2023.

²⁷In alternative versions we include up to ten newspapers and observe a high correlation and many overlapping articles, so we retain the three-paper baseline.

²⁸UK based CPU can be found in Berestycki et al., 2022

B.4 Discussion On The Risk of False Positives And False Negatives

As pointed out by the OECD (2022), it is important to be aware that when constructing an index based on article counts from newspapers, two types of errors may arise. These are either false negatives (relevant articles might be left out) and false positives (irrelevant articles are included). Out of these two error types the latter is the most important to minimize. Regarding the former, it is reasonable that in line with the climate policy index created by the OECD (2022), we suspect that our search algorithm only captures a certain number of all relevant articles for Norway. Just as the OECD, we are not able to verify whether false negatives are even distributed, but for the results to be biased, the ratio of relevant articles picked up by our search algorithm must fluctuate with macroeconomic trends and conditions. As mentioned in the OECD working paper (2022), this is unlikely, which means that our climate policy uncertainty indicator could at worst be seen as a decent proxy for the total number of articles about climate policy uncertainty in Norway. When considering false positives, we want to avoid including irrelevant articles that might use search terms such as, but not limited to, "klima" or "miljø" to describe other themes not tied to uncertainty related to climate policies. Some examples might be articles talking about "arbeidsmiljø" (meaning work environment) or "politisk klima" (meaning political climate). To reduce and try to minimize the amount of false positives picked up by the search algorithm we carefully read through hundreds of randomly selected articles across the time period. From this we reiterate the search algorithm recursively, excluding words that tend to be linked to a large portion of the irrelevant articles picked up. In our case, one such example is the term "miljø", which is a similar issue faced by the OECD (2022), but in contrast to the OECD we have not included alternative terms²⁹ due to the lack of sufficiently good Norwegian alternatives. An other example is the translation of the English search term "hydro", which in Norwegian translates to "vann". This is too vague

²⁹See Berestycki et al. (2022) Annex A page 49 for examples of alternative terms they include regarding the English translation of "miljø", which is "environment".

for the Norwegian search algorithm, thus we had to be more specific using terms such as "vannkraft" and "vanmagasin"³⁰.

C Appendix C: Raw counts and ratios for newspaper articles

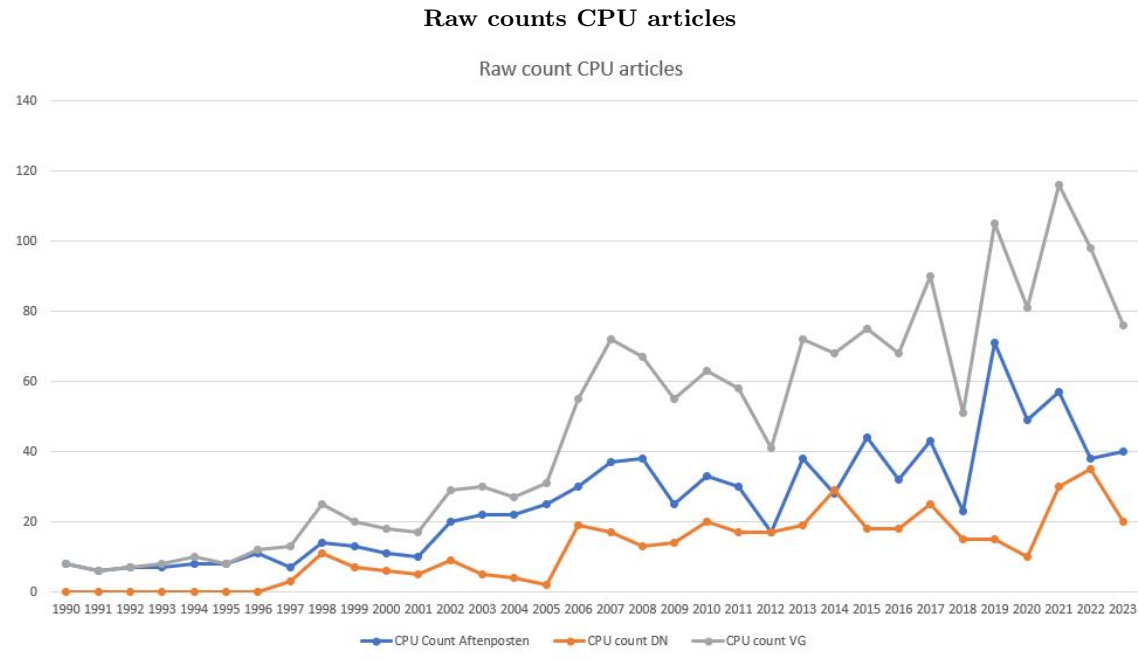


Figure C.1. Raw counts for number of climate policy uncertainty articles over time published in Aftenposten, Dagens Næringsliv, and VG. The counts are retrieved from Retriever’s Atekst database. Note: Will update graphs from excel graphs to R-generated graphs

³⁰These two translates into "hydro power" and "hydro reservoir".

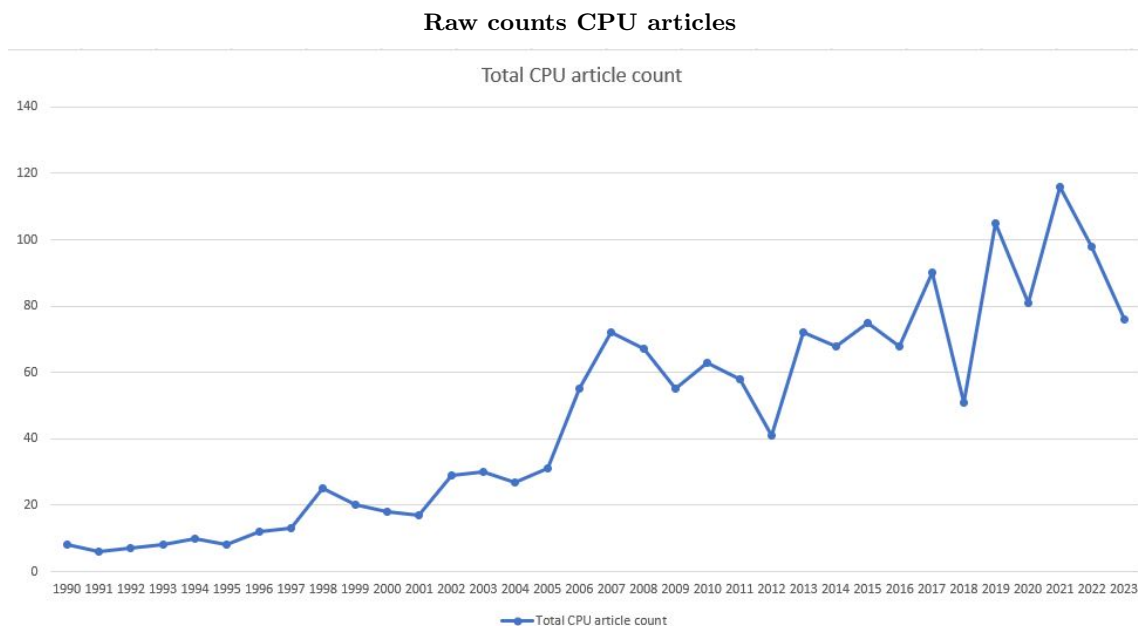


Figure C.2. Raw counts for total number of climate policy uncertainty articles over time published across the three newspapers Aftenposten, Dagens Næringsliv and VG. . The counts are retrieved from Retriever’s Atekst database. Note: Will update graphs from excel graphs to R-generated graphs

Raw counts total number of newspaper articles per newspaper

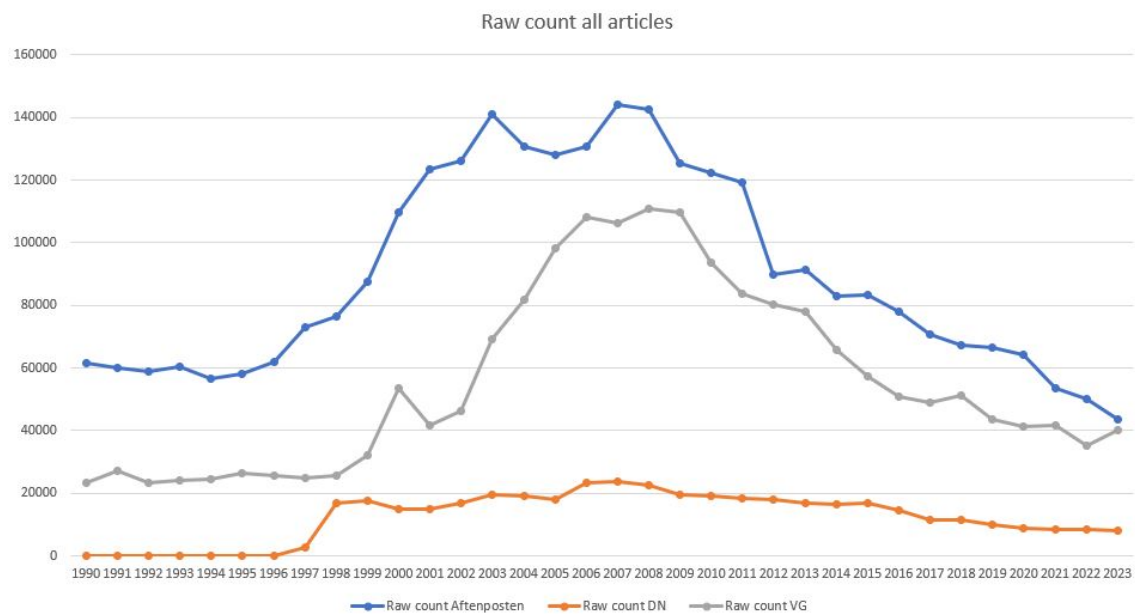


Figure C.3. Raw counts for total number of newspaper articles over time published in Aftenposten, Dagens Næringsliv, and VG. Data graphics retrieved from the Retriever search engine's analytics tool. Note: Will change excel graphs to R-generated graphs

Raw counts sum of total number of newspaper articles

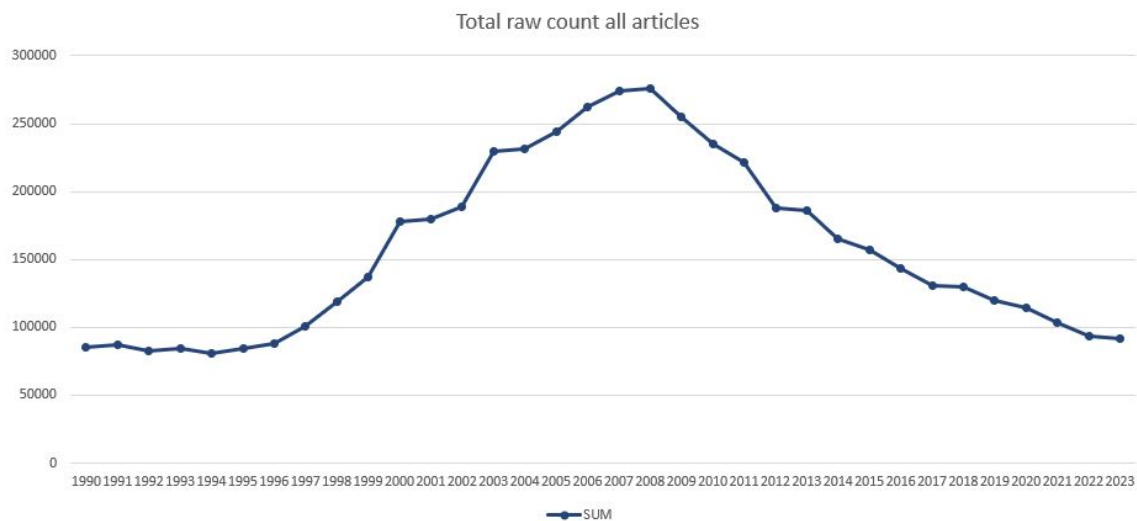


Figure C.4. Raw counts for the sum total number of newspaper articles over time published in Aftenposten, Dagens Næringsliv, and VG. Data graphics retrieved from the Retriever search engine’s analytics tool. Note: Will change excel graphs to R-generated graphs

Ratio of CPU articles for each newspaper

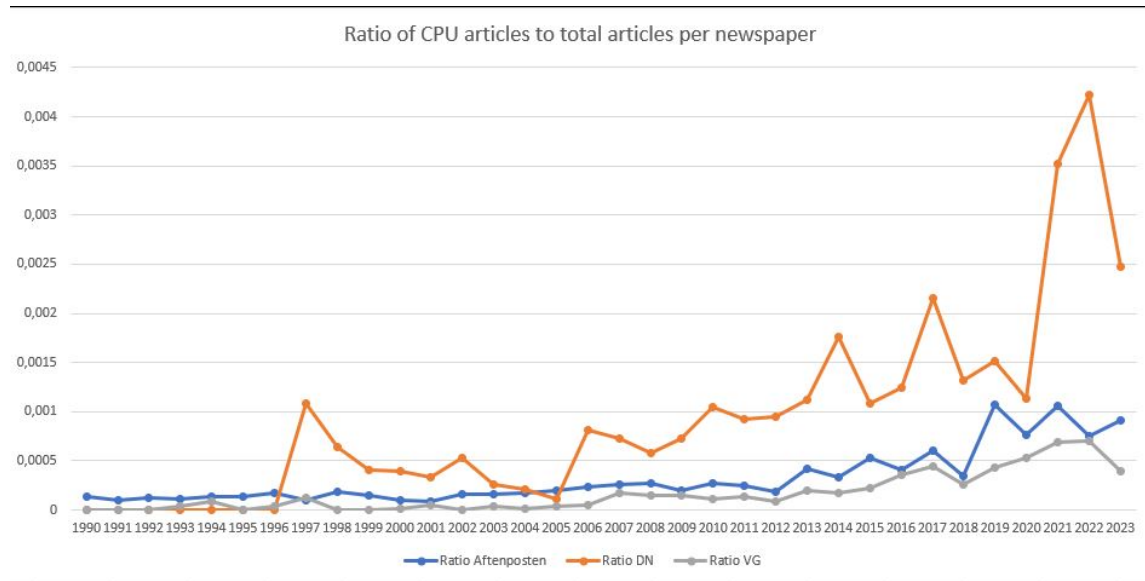


Figure C.5. Ratio's calculated by authors.

Below we also show the Norwegian Climate Policy Index at a quarterly and monthly frequency.

D Appendix D: Variations Norwegian Climate Policy Index

D.1 Quarterly Norwegian Climate Policy Uncertainty

The quarterly breakdown provides a more granulated picture than the traditional yearly climate policy uncertainty index.

Norwegian Climate Policy Uncertainty, Quarterly

Quarterly Norwegian Climate Policy Uncertainty Index




Figure D.1. Authors' own calculation. Based on the A-tekst archives. Timeline spanning from January 1995 until June 2024.

D.2 Monthly Norwegian Climate Policy Uncertainty

This gives the most granular breakdown of the Norwegian Climate Policy Uncertainty index. As compared to the climate policy uncertainty indices for the United States and the United Kingdom presented by the OECD [Berestycki et al., 2022](#), we do not have as good a monthly coverage over time.




Figure D.2. Authors’ own calculation. Based on the Retriever archives. Timeline spanning from January 1995 until June 2024.

E Appendix E: Total Annual Dollar Loss

Figure E.1 reveals a pronounced upward trend in total annual dollar losses attributable to climate-policy uncertainty across Norwegian fields: after hovering below USD 200 million through most of the 1990s and early 2000s, losses accelerate from 2013 onward, peaking near USD 700 million in 2021 before easing slightly in 2022. The pattern underscores how episodic surges in regulatory debate—most notably the post-Paris transition discussions and

the 2021-2022 energy-price turmoil—translate into markedly higher costs in the Norwegian oil industry.




Figure E.1. Total Annual Dollar Loss summed across all 55 fields in our final dataset. Numbers are in Million USD / Year

Below is the table corresponding with Figure E.1 with the yearly sum of dollar losses across all fields in our dataset.

Table E.1. Total Annual Dollar Losses, 1995–2022

Year	Total Loss, USD (M\$)	Year	Total Loss, USD (M\$)
1995	105.036	2009	186.975
1996	76.912	2010	215.571
1997	249.290	2011	176.848
1998	226.502	2012	121.187
1999	110.843	2013	272.340
2000	0.000	2014	310.344
2001	2.958	2015	351.590
2002	176.502	2016	397.806
2003	36.715	2017	593.974
2004	2.664	2018	265.777
2005	28.353	2019	591.918
2006	184.651	2020	462.636
2007	262.038	2021	803.843
2008	214.645	2022	704.142

Note: Total annual new dollar losses are expressed in million USD per year.

Data cover the period 1995–2022.

F Appendix F: Yearly average $\Delta\lambda^{CPU}$




Figure F.1. Total Annual Dollar Loss summed across all 55 fields in our final dataset. Numbers are in million USD / Year

G Appendix G: Welfare Summary Statistics, 2015-2022

G.1 Summary statistics for discovered oil in the 2015-2022 period

Table G.1. Yearly Summary Statistics, Average Across Fields, 2015–2022

	Unit of Account	Mean	Median	SD	Min	Max
$\Delta\lambda^{CPU}$	USD/BBL	1.336	1.282	0.613	0.667	2.240
Uncertainty Cost	Million USD / Year	11.017	11.201	4.251	5.778	17.475
Emissions	Million Ton CO ₂ /Year	0.292	0.295	0.021	0.251	0.318
Shadow Carbon Tax	USD/Tonne CO ₂	38.385	37.243	16.696	19.864	62.905

Note: Summary statistics for the post-Paris Agreement period (2015–2022), averaged across fields.

Table G.2. Yearly Summary Statistics, Total Sum Across Fields, 2015–2022

	Unit of Account	Mean	Median	SD	Min	Max
Fiscal Loss	Million USD / Year	521.461	527.277	183.737	265.777	803.843
Emissions	Million Ton CO ₂ / Year	15.116	15.242	1.538	12.301	16.769

Note: These totals are summed over all fields for each year in the post-Paris Agreement period (2015–2022). “Fiscal Loss” is the annual government revenue loss due to climate-policy uncertainty; “Emissions” is the total annual CO₂ emissions from these fields.

G.2 Summary statistics for undiscovered oil in the 2015-2022 period

Table G.3. Field-level shadow price of undiscovered oil, 2015–2022

	Unit of Account	Mean	Median	SD	Min	Max
μ_{\min}	USD/BBL	6.316	3.396	16.886	-24.488	42.995
μ	USD/BBL	4.988	2.525	16.511	-25.761	40.824
$\Delta\mu$	USD/BBL	1.329	1.273	0.599	0.667	2.240

Note: Summary statistics. $\Delta\mu^{CPU} = \mu_{\min}^{CPU} - \mu^{CPU}$. Data are based on 50 fields on the Norwegian Continental Shelf over the period 2015–2022.

Table G.4. Yearly Summary Statistics, Average Across Fields, 2015–2022

	Unit of Account	Mean	Median	SD	Min	Max
$\Delta\mu$	USD/BBL	1.336	1.282	0.613	0.667	2.240
Uncertainty Cost	Million USD / Year	1.467	1.234	0.856	0.552	3.315
Emissions	Million Ton CO ₂ /Year	0.039	0.043	0.011	0.022	0.051
Shadow Carbon Tax	USD/Tonne CO ₂	38.757	33.516	21.704	17.078	77.447

Note: Summary statistics for the period 2015–2022, averaged across fields.

Table G.5. Yearly Summary Statistics, Total Sum Across Fields, 2015–2022

	Unit of Account	Mean	Median	SD	Min	Max
Fiscal Loss	Million USD / Year	20.183	19.617	13.618	3.865	49.732
Emissions	Million Ton CO ₂ /Year	0.503	0.407	0.267	0.213	0.904

Note: Summary statistics for the post-Paris Agreement period. These values are the total sum across all fields each year: Fiscal Loss is the aggregate USD loss, and Emissions is the aggregate CO₂ emissions.

H Appendix H: Cost Function Parameters

The reduced-form parameters of the cost function presented in Equation (10) have the following structural interpretation:

$$\begin{aligned}
 \theta_{0t} &= \delta_0 I_t + \zeta_0 - \frac{\chi_0 + \chi_{3t}}{\chi_1} \delta_0 \alpha_4 + \alpha_0 \delta_0 - \frac{\alpha_1^2 \delta_0}{4\alpha_3} - \frac{\alpha_1 \alpha_5 \delta_0}{2\alpha_3} \\
 \theta_{1t} &= \delta_1 I_t + \zeta_1 - \delta_1 \alpha_4 \frac{\chi_{3t}}{\chi_1} + \alpha_0 \delta_1 - \frac{\alpha_1^2 \delta_1}{4\alpha_3} - \frac{\alpha_1 \alpha_5 \delta_1}{2\alpha_3} \\
 \theta_2 &= \zeta_2 + \delta_2 \\
 \theta_3 &= \zeta_3 \\
 \theta_4 &= \frac{\delta_0 \alpha_4}{\chi_1} \\
 \theta_5 &= \frac{\delta_1 \alpha_4}{\chi_1}
 \end{aligned}$$

where $I_t = \alpha_2 \mathbb{E}[Reg_t] + \left(\alpha_4 - \frac{\alpha_5^2}{4\alpha_3} \right) \mathbb{E}[Reg_t]^2$ is a measure of expected regulatory *intensity*.



NHH



NORGES HANDELSHØYSKOLE
Norwegian School of Economics

Helleveien 30
NO-5045 Bergen
Norway

T +47 55 95 90 00
E nhh.postmottak@nhh.no
W www.nhh.no

